

Online Resource

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

Online Resource 1 to the article of the same title

Andreas Peters, Rosenheim Technical University of Applied Sciences (TH Rosenheim), Rosenheim, Germany

Corresponding author: a.peters81@icloud.com · Journal of Industrial and Business Economics

Contents: A, estimation of coefficients of variation from bracket data; B, piecewise-deterministic Markov framework and sequencing condition; C, mixture-variance derivation; D, additional sensitivity results; E, extended discussion; F, reproduction package; G, effective term at risk, with proofs, bounds and computations; H, supplementary tables and figures referenced in the main text. References to sections, tables, figures and equations without a letter prefix point to the main text.

A ESTIMATION OF COEFFICIENTS OF VARIATION FROM BRACKET DATA

A.1 Overview

Neither the notified-body survey nor the manufacturer surveys publish durations of individual applications. They report the proportion of applications, or of notified bodies, falling into predefined time intervals. From such proportions this section fits parametric service-time distributions by maximum likelihood and derives the coefficient of variation (CV) for three pathways, namely certification under the MDR, certification under the predecessor MDD, and US 510(k) clearance. Interval estimates are reported as model-implied sensitivity intervals at an assumed effective sample size and not as conventional confidence intervals, because aggregate proportions support no standard sampling-distribution argument. Point estimates do not depend on the assumed effective sample size N ; what N governs is the width of the intervals.

Where the dispersion arises matters for interpreting the estimates. In the twentieth notified-body survey, 65 percent of files moved from application to a signed written agreement within two months, whereas the reported time from agreement to a new combined quality-management and product certificate spread from 6 to 12 months (8 percent of reporting bodies) to more than 24 months (10 percent) (Austrian National Public Health Institute et al., 2026b). Much of the variance that the CV measures therefore accumulates after contracting, in phases that the manufacturer cannot fully anticipate and that the body does not fully control.

A.2 Data sets and bracket definitions

Four records enter the estimation, and they differ in resolution.

A.2.1 US file-level records

For every premarket notification and every premarket approval application the FDA publishes the date of receipt, the date of decision and the decision code (Food and Drug Administration, 2026a,b). Clearance time is the difference between the two dates in calendar days, converted to months at 30.4375 days. Decisions are assigned to the federal fiscal year ending on 30 September. Three populations are used: substantially equivalent 510(k) decisions of fiscal years 2021 to 2025 ($N = 15,706$), the subset of traditional submissions without third-party review ($N = 12,795$), granted De Novo requests of the same years ($N = 173$), and original premarket approvals

of fiscal years 2016 to 2025 ($N = 311$), where the longer window is needed because original approvals are few. Supplements are excluded throughout, as are withdrawn and not-substantially-equivalent decisions. The premarket notification file also names the applicant, the submission type, whether an accredited third party reviewed the submission and the review advisory committee; applicant names are normalised by removing punctuation and corporate suffixes before the number of decisions per applicant in the window is counted, which yields 6,146 applicants behind 15,702 decisions. Clearance time here is total calendar time, so it includes the periods in which a submission rests with the applicant after a request for additional information. The European counterpart is not symmetric. MedTech Europe reports that more than half of the time from application to a quality-management certificate falls outside the review phase, in pre-review and certificate issuance (MedTech Europe, 2024), and those periods lie inside the interval measured here, whereas the period before a written agreement is signed lies outside it. Both measures capture waiting that neither authority controls, which is what a patent clock registers, but they draw the boundary at different points. A patent clock consumes exactly that quantity, whereas MDUFA goals count only agency time, which is why the main text reports the two separately.

A.2.2 European notified-body survey

Typical time from a signed written agreement to a new certificate is what the twentieth survey reports for each responding body, in four brackets. Among the 39 bodies that had issued combined quality-management and product certificates, 8 percent reported 6 to 12 months, 51 percent 13 to 18 months, 31 percent 19 to 24 months and 10 percent more than 24 months (Austrian National Public Health Institute et al., 2026b). The unit of observation is the body, not the file, and the last bracket is open, so this record identifies the distribution of body-level practice within bracket width and identifies nothing above 24 months. Every bound derived from it in Section G is therefore a bound on typical durations across bodies. A distribution across dossiers pools within-body with between-body variation and is wider by construction, so the bounds understate file-level dispersion and cannot be read as bounds on what an individual applicant faces.

Table A.1 sets out the brackets and shares. Combined certificates are the route that Class III devices and most implantable Class IIb devices must take, which is why this

Table A.1 European brackets from the twentieth notified-body survey

k	Duration	a_k (months)	b_k (months)	$\hat{p}_k$
1	6–12 months	6	12	0.08
2	13–18 months	13	18	0.51
3	19–24 months	19	24	0.31
4	> 24 months	24	∞	0.10

Note: Shares of the 39 notified bodies that had issued combined certificates, by the typical duration they reported (Austrian National Public Health Institute et al., 2026b). For estimation the brackets are treated as contiguous at 6, 12.5, 18.5 and 24.5 months.

bracket set rather than the quality-management set serves as the European reference. Class-specific distributions cannot be separated, because the survey does not disaggregate by device class; the article therefore reports one European distribution and treats class composition as one of the sources of its spread.

A.2.3 Manufacturer survey by firm size

MedTech Europe’s 2024 survey reports the shares of respondents in four brackets, 0 to 3, 4 to 8, 9 to 15 and more than 15 months, separately for SMEs and large manufacturers and separately for the pre-review, review and certificate-issuance phases of quality-management and technical-documentation assessments (MedTech Europe, 2024, Fig. 8). For the review phase of technical documentation assessments, 28 manufacturers responded; for quality-management assessments, 39 to 43. The same source reports average total durations of 19.5 months for quality-management and 21.8 months for technical-documentation certification. No other public figures distinguish firm sizes, and these measure phases rather than whole procedures.

A.2.4 Patent term restoration records

Federal Register notices determining the regulatory review period under 35 U.S.C. §156 state, for each device, the date of the investigational device exemption, the date on which the premarket approval application was initially submitted, the date of approval and the patent for which restoration is sought (Federal Register, 2026). Seventy notices published between 9 February 2018 and 26 February 2026 concern devices, and 39 of them name a patent whose filing date can be retrieved from the patent record. Such a pair does not establish the interval D between filing and submission, since the term-bearing filing is the earliest application in a patent’s chain of benefit and the notices do not identify it; the record enters the article through the phase durations and through the existence of a restoration mechanism, and no absolute effective term is derived from it. Because restoration is available only for devices that went through premarket approval, the sample is one of Class III devices with a patent worth restoring.

A.3 Distribution families

Three candidate families are fitted to each pathway. Let $\tau > 0$ denote processing time.

Log-normal. With parameters $(\mu_{\ln}, \sigma_{\ln})$,

$$f(\tau) = \frac{1}{\tau \sigma_{\ln} \sqrt{2\pi}} \exp\left(-\frac{(\ln \tau - \mu_{\ln})^2}{2\sigma_{\ln}^2}\right), \quad (\text{A.1})$$

$$\text{CV} = \sqrt{\exp(\sigma_{\ln}^2) - 1}.$$

Log-normal service times are a standard assumption in queueing applications (Kleinrock, 1975) and arise when duration is the product of several independent multiplicative factors, a plausible description of certification in which documentation review, clinical evaluation, audit and resubmission cycles combine. Two caveats apply. The upper tail of MDR durations is driven in part by resubmission after incomplete applications, which is more characteristic of a mixture or of a Weibull distribution with shape below one than of a pure log-normal. Survey brackets also end with an open interval above 24 months, so the data cannot fully separate a log-normal tail from heavier-tailed alternatives.

Weibull. With shape k_W and scale λ_W ,

$$f(\tau) = \frac{k_W}{\lambda_W} \left(\frac{\tau}{\lambda_W}\right)^{k_W-1} \exp\left[-\left(\frac{\tau}{\lambda_W}\right)^{k_W}\right], \quad (\text{A.2})$$

$$\text{CV} = \sqrt{\frac{\Gamma(1+2/k_W)}{\Gamma(1+1/k_W)^2} - 1}.$$

A shape above one implies an increasing hazard, so that longer-pending applications become more likely to be completed; a shape below one would indicate applications that get stuck.

Gamma. With shape α_G and rate β_G ,

$$f(\tau) = \frac{\beta_G^{\alpha_G}}{\Gamma(\alpha_G)} \tau^{\alpha_G-1} e^{-\beta_G \tau}, \quad \text{CV} = 1/\sqrt{\alpha_G}. \quad (\text{A.3})$$

Gamma distributions arise when processing is the sum of sequential stages with exponential durations.

A.4 Maximum likelihood from bracket data

Given K brackets $[a_k, b_k)$ with observed proportions $\hat{p}_k$, the bracket probabilities under a distribution $F(\cdot; \theta)$ are $\pi_k(\theta) = F(b_k; \theta) - F(a_k; \theta)$, and the multinomial log-likelihood at effective sample size N is

$$\ell(\theta) = \sum_{k=1}^K N \hat{p}_k \ln \pi_k(\theta). \quad (\text{A.4})$$

In the reproduction script, ℓ is maximised with the Nelder–Mead algorithm; a bounded L-BFGS-B run started from a perturbed point returns the same CV to three decimals on every pathway and family (Section F).

Sample size differs sharply between the two bracket sets. European brackets are reported by 39 notified bodies, so $N = 39$ throughout for that record, while the 510(k) brackets are formed from 12,795 file-level decisions, where sampling error is negligible and the discrepancy between fit and truth in Table A.2 is misspecification rather than noise. Intervals follow from the inverse observed Fisher information $I(\hat{\theta})^{-1}$ at the assumed N and the delta method for $g(\theta) = \text{CV}$:

$$\widehat{\text{CV}} \pm 1.96 \sqrt{\nabla g(\hat{\theta})^\top I(\hat{\theta})^{-1} \nabla g(\hat{\theta})}. \quad (\text{A.5})$$

Table A.2 Bracket fits against file-level truth, and fits to the European survey brackets

Data	Family	Mean	CV	P50	P90	$P(\tau > 24)$
FDA 510(k), traditional, FY2021–2025	file level (truth)	6.22	0.695	5.16	11.26	0.007
	log-normal on brackets	6.36	0.803	4.96	12.26	0.013
	Weibull on brackets	6.01	0.670	5.24	11.54	0.001
	gamma on brackets	6.09	0.690	5.15	11.70	0.003
EU MDR, typical time per body	log-normal	18.17	0.255	17.61	24.28	0.108
	Weibull	18.06	0.255	18.23	23.89	0.095
	gamma	18.14	0.247	17.77	24.08	0.102

Note: Months. The 510(k) brackets are 0–3, 3–4, 4–5, 5–7, 7–12 and above 12 months, with shares computed from the file-level records; the European brackets are those of the twentieth notified-body survey (6–12, 13–18, 19–24 and above 24 months, shares 8, 51, 31 and 10 percent among 39 bodies). Fits are by maximum likelihood on the grouped shares. European values describe typical durations per body, not the distribution across files. Reproduced by the script of Section F.

Table A.3 Fit comparison on the European survey brackets

Measure	Log-normal	Weibull	Gamma
Cross-entropy $\hat{H}$	1.141	1.179	1.145
Max. cumulative deviation	0.012	0.067	0.026
$\Delta\hat{H}$ vs. log-normal	0	0.038	0.004
ΔAIC at $N = 39$	0	3.0	0.3
Fitted $P(\tau > 24)$ (reported 0.10)	0.108	0.095	0.102
Fitted CV	0.255	0.255	0.247

A.5 Results

Before the method is applied to European data, it can be checked where file-level records exist. The FDA publishes the receipt and decision date of every premarket notification, so the true distribution of clearance time is observable and can be compared with what a bracket fit to the same data recovers. Grouping the 12,795 traditional 510(k) clearances of fiscal years 2021 to 2025 without third-party review into the six brackets of Table A.2 and fitting the three families to the resulting shares reproduces the mean within 2.3 percent for the log-normal family and within 3.3 percent for the other two. Dispersion is recovered less well. Against a true coefficient of variation of 0.695, the log-normal fit gives 0.803, an overstatement of 16 percent, while Weibull and gamma understate it slightly. All three families understate the far tail, giving between 0.1 and 1.3 percent above 24 months against a true 0.7 percent. Bracket fits therefore locate the centre of a distribution well and its tail only approximately, which is the reason the main text reports distribution-free bounds alongside every fitted tail quantity.

A.6 Goodness of fit

Table A.3 compares the three families on the European survey brackets. All three have two parameters and are fitted to the same four shares, so the difference in AIC between any two reduces to $\Delta\text{AIC} = 2N\Delta\hat{H}$, where $\hat{H} = -\sum_k \hat{p}_k \ln \pi_k(\hat{\theta})$ is the per-observation cross-entropy and N the number of reporting bodies. Cross-entropy and the maximum cumulative-bracket deviation do not depend on N .

Whatever the number of bodies behind the shares, the ordering stays the same. The log-normal family attains the lowest cross-entropy and the smallest cumulative deviation, gamma follows closely, and Weibull fits worst because it places too little mass above 24 months. At $N = 39$ the differences are small enough that the families are barely separable, which matters less here than it might, since all three agree on the

coefficient of variation to within about three percent. Narrow support explains why. With 82 percent of the bodies in two adjacent brackets, the fitted centre is well determined and the tail is not. Additional spread above the last reported boundary cannot be excluded, and the resubmission mechanism documented in the surveys (Huusko et al., 2023; Carl and Hochmann, 2024) gives a reason to expect it.

A.7 Replication steps

For each pathway and family, the procedure initialises the parameters, evaluates the bracket probabilities, maximises equation (A.4) at the declared N , computes the observed Fisher information numerically at the optimum, and derives the CV with its delta-method interval from equation (A.5). The script of Section F prints every quantity in Tables A.2 and A.3 for the reproduced rows.

B PIECEWISE-DETERMINISTIC MARKOV FRAMEWORK AND THE SEQUENCING CONDITION

B.1 Why a piecewise-deterministic process with review age

Regulatory status in a jurisdiction is discrete, since a dossier is either pending or cleared. A diffusion would impose continuous fluctuation on what is mechanically a binary event. Piecewise-deterministic Markov processes (PDMPs), whose state evolves deterministically between random jumps, capture that structure directly (Davis, 1993). Constant clearance intensity would, however, imply exponential waiting times with a CV of one, whereas the estimated CVs lie between about 0.25 and 0.45. The state is therefore augmented with the review age of each pending dossier, and clearance occurs with an age-dependent hazard that can be calibrated to the fitted Weibull, log-normal or gamma distribution.

B.2 State, hazard, controls and payoff

At time t the state is $(T(t), x(t), z(t)) \in [0, T^{\max}] \times \{0, 1\}^J \times \mathbb{R}_+^J$, where T is residual statutory term, x_j records clearance status in jurisdiction j and z_j the review age of a pending dossier. Residual term decays at unit rate, $\dot{T}(t) = -1$. Clearance status is absorbing; x_j starts at zero while the dossier is pending and switches to one at clearance. Review age advances at unit rate while a submission is active and pending, $\dot{z}_j(t) = \mathbf{1}\{x_j(t) = 0, a_j(t) = 1\}$, and is reset to zero at the clearance jump. Conditional on the state and control, clearance in $[t, t + dt]$ occurs with probability $h_j(z_j, a) dt + o(dt)$.

For a Weibull fit with shape k_j and scale η_j the hazard is

$$h_j(z) = \frac{k_j}{\eta_j} \left(\frac{z}{\eta_j} \right)^{k_j-1}, \quad z \geq 0, \quad (\text{B.1})$$

with $E[\tau_j] = \eta_j \Gamma(1 + 1/k_j)$. A shape above one corresponds to a CV below one; at the baseline CV of 0.35 the implied shape is about 3.1, so the hazard increases with review age, consistent with issuance accelerating once substantive review nears completion.

Admissible controls $a \in A(x) \subset \{0, 1\}^J$ indicate whether a submission is active in each jurisdiction. A submission can be activated at most once, remains active until clearance, and is inactive after clearance. The instantaneous payoff is

$$r(T, x, a) = \sum_{j=1}^J \mathbf{1}\{x_j = 1\} \pi_j(T) - \sum_{j: x_j=0} c_j(a_j), \quad (\text{B.2})$$

where $\pi_j(T) \geq 0$ is the revenue flow in jurisdiction j while term remains, parameterised as $\pi_j(T) = \bar{\pi}_j T / T^{\max}$, and c_j is the flow cost of an active submission.

B.3 Value function and dynamic-programming equation

With discount rate $\delta > 0$ the firm maximises

$$V(T, x, z) = \sup_{a(\cdot)} E_{T, x, z} \left[\int_0^\infty e^{-\delta s} r(T(s), x(s), a(s)) ds \right], \quad (\text{B.3})$$

where $E_{T, x, z}$ denotes expectation conditional on the initial state. Formally, the value function satisfies the Hamilton–Jacobi–Bellman equation with jump terms

$$\begin{aligned} \delta V = \sup_{a \in A(x)} \left\{ r(T, x, a) - \partial_T V + \sum_{j \in P(a)} \partial_{z_j} V \right. \\ \left. + \sum_{j: x_j=0} h_j(z_j, a) \Delta_j V \right\}, \end{aligned} \quad (\text{B.4})$$

where $P(a) = \{j : x_j = 0, a_j = 1\}$ collects pending jurisdictions with an active submission and $\Delta_j V \equiv V(T, x^{j \rightarrow 1}, z^{j \rightarrow 0}) - V(T, x, z)$ is the value gain at clearance, with $x^{j \rightarrow 1}$ flipping component j to one and $z^{j \rightarrow 0}$ resetting its review age. The left-hand side is the required return on current value. On the right, r is the running payoff, $-\partial_T V$ captures the decay of term, the age terms capture the progression of review, and the non-local jump terms add the value gain at clearance weighted by the hazard. Boundary conditions are $V(0, x, z) = 0$ and, once every jurisdiction is cleared, $V(T, \mathbf{1}, z) = \sum_j \int_0^T e^{-\delta s} \pi_j(T-s) ds$.

B.4 Regularity and well-posedness

Assumption 1 (Lipschitz payoff). $r(T, x, a)$ is Lipschitz in T uniformly in (x, a) . Under $\pi_j(T) = \bar{\pi}_j T / T^{\max}$ the constant is $L_r = \sum_j \bar{\pi}_j / T^{\max}$.

Assumption 2 (Locally bounded hazard). For each j , h_j is continuous in z and bounded on $[0, \bar{z}]$ for every finite $\bar{z}$, which holds for the Weibull hazard with finite shape; clearance jumps therefore occur with finite intensity on bounded horizons.

Assumption 3 (Bounded payoff). $|r(T, x, a)| \leq M_r < \infty$, since revenues are bounded by $\sum_j \bar{\pi}_j$ and submission costs by their finite maximum.

Assumption 4 (Compact controls). $A(x)$ is a finite, and hence compact, subset of $\{0, 1\}^J$.

Given Assumption 3, the value function is well defined and bounded, $\sup |V| \leq M_r / \delta$, and Assumptions 1, 2 and 4 are the kind of conditions under which the control theory of PDMPs characterises value functions through dynamic-programming equations of the form (B.4) (Davis, 1993; Fleming and Soner, 2006). A complete existence and uniqueness proof for the age-augmented case is not attempted here, and none of the main-text results depends on one, because R1 is a reconstruction, P1 uses the first-order expansion below and P2 is a variance identity. Numerical treatment of Hamilton–Jacobi–Bellman systems arising in innovation dynamics is discussed by Bañas et al. (2022).

Remark 1. The formulation contains no diffusion in the discrete component, and no second derivative with respect to clearance status appears. Dispersion of clearance time enters through the shape of the age-dependent hazard, not through a noise term.

B.5 Two jurisdictions and the sequencing condition

With $J = 2$ and $j \in \{EU, US\}$, the firm compares three strategies, namely parallel entry, US-first entry with the EU submission made after US clearance, and entry in the United States alone. Let time run from the date at which submission is possible, let revenue in each market be a constant flow π from clearance until expiry at T^{res} , and let clearance times be independent. The value at time zero of an EU submission made at time t with clearance distribution F is

$$\begin{aligned} V_t(F) = E_F \left[\int_{\min(t+\tau, T^{\text{res}})}^{T^{\text{res}}} \pi e^{-\delta u} du \right] \\ - e^{-\delta t} (K + c\pi E_F[\lfloor \tau/R \rfloor]), \end{aligned} \quad (\text{B.5})$$

with the refinancing cost booked at submission. Value in the United States is the same under both entry strategies, so the comparison reduces to $V_0(F_0)$ against $E[V_{\tau_{US}}(F_1)]$, while entry in the United States alone adds zero in the EU.

Derivation of Proposition 1 of the main text. Without expiry truncation, $V_t(F) = e^{-\delta t} V(F)$ with $V(F)$ as in equation (5) of the main text, and independence gives $E[V_{\tau_{US}}(F_1)] = E[e^{-\delta \tau_{US}}] V(F_1)$, which is the condition of the proposition. If $F_1 = F_0$ and $V(F_0) > 0$, then $E[e^{-\delta \tau_{US}}] V(F_0) < V(F_0)$ because $E[e^{-\delta \tau_{US}}] < 1$, so parallel entry is strictly better. When $V(F_0) < 0$, US-first entry loses less than parallel entry but still loses against entry in the United States alone. With truncation, a later submission leaves less term, so V_t falls faster than $e^{-\delta t} V$ and both orderings are reinforced. $\square$

Table B.1 evaluates the comparison with truncation over the documented range of every input. Entries are the reduction of the mean European clearance time, at constant coefficient of variation, that makes US-first entry exactly as good as parallel entry; a firm expecting a larger reduction from completing its file during the US review should sequence. Three regularities stand out. The required reduction falls as the submission cost rises, because a costly European submission is better deferred, and it falls from 5.31 to 3.12 months behind a 510(k) review for the smallest size class once the full first-year certification cost is counted rather than notified-body fees alone. Refinancing has a similar effect, lowering the requirement by about one

Table B.1 Reduction of mean European clearance time required for US-first entry, by cost basis, cost of capital and refinancing cost (months)

Cost basis, persons employed	$\delta = 0.0754$			$\delta = 0.0957$			$\delta = 0.1057$		
	$c = 0$	0.55	0.90	$c = 0$	0.55	0.90	$c = 0$	0.55	0.90
<i>US route: 510(k) clearance</i>									
Fees, 10–19	5.35	4.36	3.80	5.31	4.28	3.72	5.29	4.25	3.68
Fees, 20–49	5.43	4.43	3.87	5.41	4.38	3.81	5.40	4.35	3.78
Fees, 50–249	5.47	4.46	3.91	5.46	4.42	3.85	5.46	4.40	3.82
Total, 10–19	3.68	2.86	2.46	3.12	2.36	2.01	2.83	2.11	1.78
Total, 20–49	4.84	3.89	3.37	4.64	3.68	3.17	4.54	3.57	3.07
<i>US route: premarket approval</i>									
Fees, 10–19	13.90	12.79	12.07	13.70	12.55	11.80	13.60	12.43	11.67
Fees, 20–49	14.11	12.99	12.28	13.96	12.81	12.07	13.89	12.72	11.97
Fees, 50–249	14.20	13.08	12.37	14.08	12.93	12.19	14.02	12.85	12.10
Total, 10–19	9.63	8.48	7.74	8.16	6.96	6.21	7.41	6.19	5.45
Total, 20–49	12.61	11.48	10.76	12.03	10.86	10.11	11.75	10.55	9.78

Note: Inputs from Table 3 of the main text. K/π is the ratio of notified-body fees (313,183 euros) or total first-year certification cost (about 4.47 million euros) to net turnover per enterprise in the relevant size class; c is the cost of one further financing round in years of revenue; the financing interval is 22.9 months; valuation horizon eleven years with expiry truncation, to which the entries are insensitive (varying it between eight and twenty years moves them by at most 0.11 months). 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, at $\delta = 0.0957$ and $c = 0.55$ with fee-based costs for the 20–49 class.

to one and a half months at the higher of the two documented dilution levels. The cost of capital matters least, since across its documented range the requirement moves by less than one month in every row except the three in which a firm of 10 to 19 persons employed bears the full first-year cost behind a premarket approval, where it reaches 2.3 months. Dispersion of the European clock enters the comparison through the refinancing term alone, because at a constant mean a wider distribution raises the expected number of financing intervals that certification spans; holding the mean at the survey value and raising the coefficient of variation from 0.26 to 0.70, the value measured in the SME review phase, lowers the required saving by a further 0.4 to 0.9 months at $c = 0.55$. The comparison is practically independent of the interval D , because truncation rarely binds over the range of residual terms examined, from eight to twenty years. Behind a premarket approval the requirement is between 5.5 and 14.2 months, because the US review itself consumes 14.7 months on average, and no documented difference between firm types approaches that figure.

Endogenous preparation time. The completion of the European file during a US review is one instance of a general trade-off. If a firm chooses file quality q , preparation takes $D(q) = \underline{D} + \phi(q)$ with $\phi' > 0$, and clearance follows F_q with mean and dispersion falling in q , then effective term is $T^{\max} - \underline{D} - \phi(q) - \tau(q)$. Expected term rises with q only if completeness saves more queue time than it costs in preparation, while the tail can improve even when the mean hardly moves, because resubmission rounds fill the right tail (Section D.8). A higher value of completeness arises from dispersion only through frictions, since the risk-neutral value in equation (B.6) is convex in τ ; a short runway or a steep preparation cost lowers the chosen q and thereby lengthens and widens the clearance distribution, which gives the SME parameters of Proposition 2 of the main text a possible behavioural foundation. None of these comparative statics is calibrated here.

Omitted channels. US revenue that finances the European wait would reduce the need for a further financing round under US-first entry and strengthen its case. Information from US

Table B.2 Effect of dispersion on risk-neutral expected present value, in days of equivalent delay reduction

Distribution	$\delta=0.0754$	0.0957	0.1057
EU MDR, survey fit (18.17; 0.25)	2.0	2.6	2.8
EU MDR, TDA review, SMEs (18.7; 0.70)	15.3	19.2	21.0
FDA 510(k), file level	1.8	2.3	2.5
FDA PMA originals, file level	17.3	21.3	23.2

Note: valuation horizon eleven years; discount rates are the US and European healthcare cost of capital and the European cost of equity (Damodaran, 2026). US rows integrate over the file-level records, European rows over the fitted log-normal distributions. Reproduced by the script of Section F.

clearance could shorten the European clock where clinical data from a premarket approval can be reused; for 510(k) clearances that rest on substantial equivalence without new clinical data such spillovers are less plausible. The value of waiting for information about the European queue itself is not computed.

B.6 Valuation curvature and the channels of dispersion cost

Consider a risk-neutral firm that earns a constant revenue flow π from clearance at τ until expiry at T^{res} , both measured in years from submission, and discounts at rate δ . Present value at submission is

$$\text{PV}(\tau) = \frac{\pi}{\delta} (e^{-\delta\tau} - e^{-\delta T^{\text{res}}}), \quad \text{PV}''(\tau) = \pi \delta e^{-\delta\tau} > 0, \quad (\text{B.6})$$

so value is convex in clearance time. A second-order expansion around $\mu = E[\tau]$ gives $E[\text{PV}(\tau)] - \text{PV}(\mu) \approx \frac{1}{2} \text{PV}''(\mu) \sigma^2$, and dividing by the marginal cost of delay $|\text{PV}'(\mu)| = \pi e^{-\delta\mu}$ expresses the effect as an equivalent reduction of expected delay by $\frac{1}{2} \delta \sigma^2$ years. Table B.2 reports exact values, computed over the file-level US records and over the fitted European distributions. Dispersion alone raises expected present value by at most three days on the European pathway as the survey describes it and by about three weeks where dispersion is as wide as in SME review phases or in US premarket approvals, which is still small against clearance times measured in years.

Once frictions enter, dispersion becomes costly. With a concave valuation by the owner, for instance a mean-variance certainty equivalent $E[PV] - \frac{a}{2} \text{Var}[PV]$ with $a > 0$, the relevant variance is approximately $PV'(\mu)^2 \sigma^2$. At the European cost of capital the standard deviation of present value amounts to 6.1 percent of its mean for European certification as the survey describes it, 15.6 percent for the SME review phase, 5.3 percent for a 510(k) clearance and 15.1 percent for an original premarket approval, so on this measure the European survey distribution and the 510(k) records are close, while the SME review phase and the premarket approval route are three times more dispersed. Weight in the owner's decision grows with the share of wealth the project represents, which is highest for single-product firms.

C MIXTURE-VARIANCE DERIVATION AND SME AMPLIFICATION

C.1 Statement and proof

Proposition 1 (SME amplification through mixture variance; Proposition 2 of the main text). *Let τ be the processing time of a randomly drawn application and $Z \in \{L, S\}$ the firm type, with $P(Z = L) = p$ and $P(Z = S) = 1 - p$. Conditionally, $\tau | Z = L$ has mean μ_L and variance σ_L^2 , and $\tau | Z = S$ has mean μ_S and variance σ_S^2 , with $\mu_S > \mu_L$ and $\sigma_S > \sigma_L$. Then*

$$H^{\text{port}} = \underbrace{p\sigma_L^2 + (1-p)\sigma_S^2}_{\text{within-group}} + \underbrace{p(1-p)(\mu_L - \mu_S)^2}_{\text{between-group}}, \quad (\text{C.1})$$

and $H^{\text{port}} > p\sigma_L^2 + (1-p)\sigma_S^2$ whenever $\mu_L \neq \mu_S$ and $0 < p < 1$.

Proof. By the law of total variance, $\text{Var}[\tau] = E_Z[\text{Var}(\tau | Z)] + \text{Var}_Z[E(\tau | Z)]$. In that decomposition the first term equals $p\sigma_L^2 + (1-p)\sigma_S^2$. The conditional mean $m(Z)$ takes the value μ_L with probability p and μ_S with probability $1 - p$, so

$$\begin{aligned} \text{Var}_Z[m(Z)] &= p\mu_L^2 + (1-p)\mu_S^2 - [p\mu_L + (1-p)\mu_S]^2 \\ &= p(1-p)\mu_L^2 + p(1-p)\mu_S^2 - 2p(1-p)\mu_L\mu_S \\ &= p(1-p)(\mu_L - \mu_S)^2. \end{aligned} \quad (\text{C.2})$$

Combining both terms gives equation (C.1); strict inequality follows because $p(1-p)(\mu_L - \mu_S)^2 > 0$ under the stated conditions. $\square$

C.2 Interpretation

Portfolio variance splits into a weighted average of the variance each firm type experiences and an amplification term that arises purely from differences in conditional means; the latter would exist even if both firm types faced no dispersion of their own. Pooled statistics therefore understate the dispersion that SMEs face and overstate the dispersion that large firms face.

Corollary 1. *SME-specific variance exceeds pooled portfolio variance, $\sigma_S^2 > H^{\text{port}}$, if and only if*

$$\sigma_S^2 > \sigma_L^2 + (1-p)(\mu_L - \mu_S)^2. \quad (\text{C.3})$$

Proof. Substituting (C.1) into $\sigma_S^2 > H^{\text{port}}$ and subtracting $(1-p)\sigma_S^2$ gives $p\sigma_S^2 > p\sigma_L^2 + p(1-p)(\mu_L - \mu_S)^2$; dividing by $p > 0$ yields (C.3). $\square$

Table C.1 Mixture-variance decomposition of the review phase

Group	μ (months)	σ (months)	σ^2	Share
Large firms	13.37	4.63	21.5	0.60
SMEs	18.66	13.06	170.7	0.40

Within-group: $0.60 \times 21.5 + 0.40 \times 170.7 = 81.2$

Between-group: $0.60 \times 0.40 \times (18.66 - 13.37)^2 = 6.72$

Pooled: $H^{\text{port}} = 87.9$ ($\sigma^{\text{port}} = 9.37$ months); between-group share 7.6 percent




Fig. C.1 Portfolio variance decomposition as a function of the large-firm share p at the base-case values of Table C.1. The between-group component is largest at $p = 0.5$ and vanishes at $p \in \{0, 1\}$

At the base-case values of Table C.1 the right-hand side equals $16 + 0.6 \times 36 = 37.6$ squared months, well below $\sigma_S^2 = 56.25$. The margin grows with the mean gap for fixed within-group dispersion.

C.3 Numerical examples

MedTech Europe's 2024 survey reports, for the review phase of technical documentation assessments under the MDR, the shares of respondents in four duration brackets separately for SMEs and for large manufacturers (MedTech Europe, 2024, Fig. 8). Half of the SMEs and a quarter of the large manufacturers reported more than 15 months, 30 and 56 percent reported 9 to 15 months, and 20 and 11 percent reported 4 to 8 months. Log-normal fits to these shares give the moments of Table C.1, and the SME share of the certificates required, 40 percent (MedTech Europe, 2022), fixes the weights.

Fig. C.1 plots the between-group component against the weight. An SME drawing from its conditional distribution faces a standard deviation of 13.06 months, which is 39 percent above the pooled value of 9.37 months and almost three times the large-firm value. The decomposition locates the SME burden in the within-group term rather than in the between-group term, and the reason is visible in the moments: the two means differ by 5.29 months, while the two standard deviations differ by 8.4 months. Where a mean gap dominates, reporting by firm size suffices to reveal it; where a dispersion gap dominates, the tail measures of Section G are needed.

Because the between-group component scales with the square of the mean gap and with $p(1-p)$, its share of pooled variance stays modest across every plausible weighting. Table C.2 varies the large-firm share of certificates, and the between-group share moves only between 5.7 and 8.1 percent, so the conclusion does not depend on the weight.

Table C.2 Between-group variance and its share under alternative weights

p	Within-group (months ²)	Between-group (months ²)	Pooled (months ²)	Between share (percent)
0.4	110.99	6.72	117.71	5.7
0.5	96.08	7.00	103.07	6.8
0.6	81.16	6.72	87.87	7.6
0.7	66.24	5.88	72.12	8.1

Note: p is the share of certificates attributable to large manufacturers; moments as in Table C.1. The reported shares for large manufacturers sum to 92 percent rather than 100; the fits normalise them proportionally. Assigning the missing mass entirely to the 0–3, 4–8 or above-15 bracket gives between-group shares of 4.7, 8.8 and 5.5 percent and pooled ninetieth percentiles of 29.0, 25.3 and 25.9 months. The central case $p = 0.6$ follows from the SME share of 40 percent reported by MedTech Europe (2022).

C.4 Conditions for between-group dominance

The between-group component exceeds the within-group component when

$$(\Delta\mu)^2 > \frac{p\sigma_L^2 + (1-p)\sigma_S^2}{p(1-p)}. \quad (\text{C.4})$$

At the base case the right-hand side equals $40.15/0.24 = 167.3$, which requires $\Delta\mu > 12.9$ months; under current conditions the within-group component dominates, and the between-group term is a sizeable amplifier rather than the leading component. If capacity expansion reduced within-group dispersion, for example to $\sigma_L = 2$ and $\sigma_S = 4$ months, without closing the mean gap, the threshold would fall to $(0.4 \times 4 + 0.6 \times 16)/0.24 = 46.7$, and a gap above 6.8 months would suffice for dominance. Reforms that reduce dispersion within groups while the gap persists thus raise the relative weight of the between-group term.

C.5 Connection to the policy argument

Aggregate capacity expansion, whether through additional designations or higher throughput per body, reduces within-group variance but leaves the mean gap between firm types largely unaffected where that gap reflects documentation readiness and access to bodies rather than aggregate capacity. Interventions that close the gap directly, such as pre-submission support for SMEs, designated review tracks or fee structures for first-time applicants, act on the between-group component at its source. The decomposition therefore identifies a direction of reform, not only an exposure.

C.6 Tail exposure by firm type

Variances understate how differently the two groups experience the upper tail. Under the fitted distributions, 40 percent of SME review phases exceed 18 months and 26 percent exceed the median financing interval of 22.9 months, against 15 and 4 percent for large manufacturers. At the ninetieth percentile the two groups are 14.9 months apart, 34.3 against 19.5 months, where their means differ by only 5.29 months.

Table C.3 repeats the exercise for the quality-management phase, where the reported shares differ less sharply, and adds the totals reported at firm level. Two points follow. The firm-size gap sits in the review phase rather than in the quality-management phase, and total certification times hardly differ by size, because SMEs report shorter pre-review phases (MedTech Europe, 2024). A comparison of total durations by

firm size would therefore find almost nothing, while the phase in which files are examined and deficiencies raised shows a pronounced difference in spread.

D ADDITIONAL SENSITIVITY RESULTS

D.1 Prior-experience count and within-applicant gradient

Table D.1 reports the two specifications that answer the look-ahead objection to the frequency classification of the main text. The first replaces the window count by the number of substantially equivalent decisions the same applicant had obtained strictly before the focal submission was received, computed over the whole public record from 1996 onwards. The second holds the applicant constant.

D.2 Alternative populations and periods

Table D.2 collects the alternatives that the records themselves offer for each pathway, namely different populations of decisions, different observation periods and, for the EU, the two sources that report duration at all. Mean clearance time under the MDR exceeds every US comparator in every combination, and the ordering of the three US pathways is stable across periods. Within the 510(k) route the fiscal years 2022 and 2023 stand out, with means of 6.70 and 6.92 months against 5.42 months in 2025, which is consistent with the backlog that followed the pandemic years.

D.3 Distribution families and the limits of distinguishability

Table D.3 restates the coefficient of variation by family for the two bracket sets that public records disclose, and Table D.4 reports the AIC comparison. On the European survey brackets the three families agree to within three percent, because the reported shares concentrate in two adjacent brackets and pin the centre without informing the tail. On the file-level 510(k) brackets, where the truth is known, the log-normal overstates the coefficient of variation by 16 percent and the other two families understate it by 1 to 4 percent, so the family choice matters more than the bracket resolution suggests and the European value of 0.26 is best read as a precision ceiling. Substituting Weibull or gamma estimates changes no qualitative conclusion.

On both bracket sets the log-normal attains the lowest cross-entropy, gamma follows and Weibull is a clear third, since a shape fitted to the main mass of the distribution cannot accommodate the share above 24 months. With 39 reporting bodies the differences in AIC are too small to separate the families, and none of this rules out a heavy-tailed alternative: the open last bracket could hide a longer tail, in which case a heavy-tailed family would fit equally well and imply a higher coefficient of variation. The survey data cannot decide between these readings, which is why every European tail quantity in the main text carries distribution-free bounds beside the fit. Log-normal serves as the primary family because it is preferred or indistinguishable on every disclosed pathway, because it is a conventional service-time model, and because its close agreement with the gamma indicates that the choice among well-behaved families does not drive the findings.

Table C.3 Firm-size differences by assessment phase

Phase and group	N	Mean	CV	P50	P90	$P(\tau > 18)$	$P(\tau > 22.9)$	Between share
TDA review, large manufacturers	28	13.37	0.347	12.63	19.45	0.147	0.039	
TDA review, SMEs	28	18.66	0.700	15.29	34.34	0.398	0.262	7.6 percent
QMS review, large manufacturers	39–43	10.28	0.967	7.39	20.93	0.137	0.082	
QMS review, SMEs	39–43	14.73	1.018	10.32	30.42	0.255	0.173	3.1 percent
Total QMS certification, all firms	39–43	19.50	–	–	–	–	–	
Total TDA certification, all firms	28	21.80	–	–	–	–	–	

Note: Months. Log-normal fits to the firm-size shares of MedTech Europe (2024, Fig. 8); totals as reported there. The between share uses $p = 0.6$. Reproduced by the script of Section F.

Table D.1 Backward-looking experience and within-applicant gradient

	OLS	$q = 0.50$	$q = 0.90$
<i>Prior decisions before the focal receipt, reference twenty or more</i>			
None	0.199 (0.021)	0.192	0.178
One to four	0.128 (0.021)	0.103	0.099
Five to nineteen	0.081 (0.022)	0.069	0.073
Observations	15,702	15,702	15,702
R^2	0.414		
<i>Within applicant, applicants with at least three submissions</i>			
Log running submission number, decision cohorts	-0.433 (0.021)		
Log prior decisions	-0.261 (0.036)		
Observations / applicants	9,895 / 1,453		
Decisions FY2021–2023, log order	-0.478		
Decisions FY2021–2022, log order	-0.438		
<i>Receipt cohorts, applicants with at least three submissions</i>			
	Receipt-year effects	Decision-year effects	Obs. / applicants
Received FY2018–2021, log order	-0.045 (0.021)	-0.390 (0.021)	7,518 / 1,103
Received FY2019–2022, log order	-0.016 (0.022)	-0.399 (0.021)	7,456 / 1,140
Received FY2018–2021, log reversed order	-0.011 (0.024)	0.393 (0.021)	7,518 / 1,103

Note: Dependent variable log clearance time in months. All columns include submission type, third-party review, applicant country, fiscal-year and device-area effects; the lower panel adds applicant fixed effects and drops applicant country, which they absorb. Standard errors clustered by applicant in parentheses. Right-censoring removes files received late in the window that are still open, which selects late submissions towards short durations and therefore biases the order coefficient away from zero. Drawing the sample by receipt cohort attenuates that selection. The lower panel reports both year specifications for each cohort, because they answer different questions: the decision year is the receipt date plus the duration on the left-hand side, so conditioning on it conditions on the outcome and reinstates the selection the cohort design is meant to remove. Receipt cohorts are built on all substantially equivalent decisions with a receipt date in the window and a positive duration, ranked within applicant by receipt date and then K-number, using the reconstructed applicant identity described in Section F; the register of decided files remains an incomplete receipt cohort, since applications still open or withdrawn never appear in it. The receipt-year specification of the first window is the one the main text reports.

D.4 Single-server comparison of exponential and general service

An M/M/1 queue, with exponential service and a CV of one, is a common simplification. Table D.5 compares sojourn times from equation (2) of the main text at the two traffic intensities used in the text, and Fig. D.1 shows the relation across the range. Exponential service overstates sojourn times by about 49 percent at $\rho = 0.70$ and about 66 percent at $\rho = 0.85$ relative to the general distribution fitted to the survey. The comparison concerns comparative statics; levels of W are not used anywhere in the analysis.




Fig. D.1 Expected sojourn time against traffic intensity under general service (two CV values) and exponential service, with $E[S] = 18.17$ months, the mean of the survey fit. The shaded range marks the traffic intensities used in the text; only the comparative statics are used

D.5 Sequencing thresholds against documented magnitudes

Table B.1 expresses the sequencing comparison as the reduction of mean European clearance time that would make US-first entry worthwhile. Read as a decision rule, it says that behind a 510(k) review a firm should sequence if it expects completion of the European file during the US review to save more than about two to five and a half months, and behind a premarket approval only if it expects to save more than about six months in the most favourable combination and more than a year in most. Two documented magnitudes can be held against those thresholds. Between firm types, the fitted mean of the technical-documentation review phase differs by 5.29 months (MedTech Europe, 2024). Within firms, 77 percent of respondents observed faster conformity assessment for a second technical-documentation certificate than for the first, while the same survey reports no comparable gain for repeated quality-management assessments (MedTech Europe, 2024). Both magnitudes clear the 510(k) threshold and neither clears the premarket-approval threshold, which is the asymmetry that the main text reports.

D.6 Sensitivity to the valuation horizon

Absolute effective terms are not reported in the main text, because the date from which a term-bearing patent runs is the earliest application in its chain of benefit and that date is not established here for any sample of devices. What the valuation results do require is a horizon over which revenue accrues, and they are almost invariant to it. Setting the residual term at eight, eleven, fifteen or twenty years changes the reduction of mean European clearance time required for US-

Table D.2 Clearance time under alternative populations and periods

Record	Population	N	Mean	CV	P90	$P(\tau > 22.9)$
EU MDR	Twentieth survey, typical time per body	39	18.17	0.255	24.28	0.149
	Same, distribution-free bounds	39	–	–	[19, 24]	[0.10, 0.41]
	MedTech Europe, average QMS certification	39–43	19.50	–	–	–
	MedTech Europe, average TDA certification	28	21.80	–	–	–
FDA 510(k)	All substantially equivalent, FY2021–2025	15,706	5.51	0.796	10.38	0.007
	Traditional, no third party, FY2021–2025	12,795	6.22	0.695	11.26	0.008
	Traditional, no third party, FY2019	2,311	5.40	0.620	9.36	–
	Traditional, no third party, FY2023	2,673	6.92	0.746	13.92	–
	Traditional, no third party, FY2025	2,612	5.42	0.604	8.80	–
FDA De Novo	Granted, FY2021–2025	173	12.12	0.573	21.39	0.081
FDA PMA	Original approvals, FY2016–2025	311	14.72	0.971	27.04	0.138
	Original approvals, FY2021–2025	148	16.24	1.087	28.92	–

Note: Months. Fiscal years end on 30 September and refer to the date of decision. The European rows describe typical durations per notified body and firm-reported averages; neither is a distribution across files. Reproduced by the script of Section F.

Table D.3 Coefficient of variation by distribution family

Bracket set	Log-normal	Weibull	Gamma
EU MDR, typical time per body	0.255	0.255	0.247
FDA 510(k), traditional, bracketed	0.803	0.670	0.690
FDA 510(k), traditional, file level		0.695	

Note: Fits as in Table A.2. No PMA brackets are needed, since PMA durations are observed file by file.

Table D.4 AIC comparison across distribution families on the European survey brackets

Measure	Log-normal	Weibull	Gamma
Cross-entropy $\hat{H}$	1.141	1.179	1.145
ΔAIC at $N = 39$	0	3.0	0.3
ΔAIC at $N = 100$	0	7.6	0.8
Max. cumulative deviation	0.012	0.067	0.026

Note: $\Delta\text{AIC} = 2N\Delta\hat{H}$; $N = 39$ is the number of reporting bodies, $N = 100$ a hypothetical larger sample. Differences below 2 do not separate the families.

first entry by at most 0.11 months behind a premarket approval and leaves it unchanged to two decimals behind a 510(k) review, because expiry truncation almost never binds at those horizons. Differences between pathways in term consumed do not involve the horizon at all.

D.7 Tail probabilities across the dispersion band

European dispersion is the quantity that public records leave open, so Table D.6 holds the mean at the survey value of 18.17 months and varies the coefficient of variation across the range that the other records display: 0.260 from the survey fit itself, 0.350 and 0.450 as intermediate values, and 0.700 as measured in the SME review phase. Raising dispersion

Table D.5 Sojourn times under M/G/1 and M/M/1 at two traffic intensities

Model	$\rho = 0.70$		$\rho = 0.85$	
	W (months)	Δ (%)	W (months)	Δ (%)
M/G/1, CV = 0.25 (survey fit)	40.7	+0.0	73.0	+0.0
M/G/1, CV = 0.70	49.8	+22.1	94.9	+30.0
M/M/1, CV = 1.00	60.6	+48.7	121.1	+66.0

Note: $E[S] = 18.17$ months, the mean of the survey fit; W from the Pollaczek–Khinchine relation.

Table D.6 European quantiles and tail exposure across the dispersion band

CV	P10	P50	P90	$P(\tau > 18)$	$P(\tau > 24)$	$E[(\tau - 22.9)^+]$
0.26	12.7	17.6	24.4	46.4%	11.2%	0.52
0.35	11.1	17.2	26.5	44.4%	16.2%	1.02
0.45	9.6	16.6	28.7	42.4%	19.4%	1.62
0.70	6.6	14.9	33.4	38.2%	22.5%	3.06

Note: Quantiles and shortfall in months; log-normal distribution at a fixed mean of 18.17 months; the financing interval is 22.9 months. Reproduced by the script of Section F.

lowers the median slightly and raises every upper quantile, and the effect on the quantity that matters for financing is large: expected shortfall beyond the median financing interval rises from 0.52 to 3.06 months, a factor of almost six, while the probability of exceeding 18 months barely moves, because 18 months lies close to the mean. Dispersion matters for horizons in the tail and hardly at all for horizons near the centre.

D.8 Sources of dispersion

Resubmission compounds the clock. Let each assessment round have mean m_S and coefficient of variation CV_S , and let a further round follow with probability r , independently, so that the number of rounds N is geometric with $E[N] = 1/(1-r)$ and $\text{Var}[N] = r/(1-r)^2$. For the random sum $\tau = \sum_{i=1}^N S_i$, $\text{Var}[\tau] = E[N] \text{Var}[S] + \text{Var}[N] m_S^2$, hence

$$CV_\tau^2 = (1-r)CV_S^2 + r. \quad (\text{D.1})$$

Table D.7 evaluates equation (D.1), and the FDA records give it a benchmark. Traditional 510(k) clearances show a coefficient of variation of 0.695 across files; generating that figure from single rounds with a coefficient of variation of 0.25 would require a resubmission probability of 0.45. Requests for additional information are indeed common on that route, and the agency has reported them on the first review cycle for a majority of submissions. Comparable figures for rounds of questions under the MDR are not published, which is one reason why the European coefficient of variation across files cannot be pinned down.

Congestion adds a transitional component. For an M/G/1 queue the waiting time satisfies $E[W_q] = \lambda_a E[S^2]/(2(1-\rho))$ and $E[W_q^2] = 2E[W_q]^2 + \lambda_a E[S^3]/(3(1-\rho))$, with arrival rate $\lambda_a = \rho/E[S]$, and the time in the system is $W = W_q + S$ with

Table D.7 Coefficient of variation of clearance time under resubmission

	$r = 0$	$r = 0.10$	$r = 0.25$	$r = 0.45$	$r = 0.65$
$CV_S = 0.20$	0.200	0.335	0.520	0.688	0.820
$CV_S = 0.25$	0.250	0.369	0.541	0.702	0.830

Table D.8 Mean and coefficient of variation of time in the system under M/G/1

ρ	0.30	0.50	0.60	0.70	0.85
$E[W]$ (months)	22.3	27.8	32.7	40.7	73.0
CV_W	0.438	0.565	0.635	0.712	0.842

independent components (Kleinrock, 1975). With the service distribution fitted to the twentieth survey, Table D.8 shows that the coefficient of variation of W falls with the traffic intensity towards the service-time value of 0.255. Even at a traffic intensity of 0.30 it is 0.438, so a large part of the dispersion an individual file meets would be queueing rather than assessment, and the coefficient of variation of 0.255 measured across bodies is a lower bound for what a file experiences. Levels of W carry no empirical weight, since a single server stands in for some fifty bodies.

Dispersion between bodies can be bounded from the twentieth survey. Among the 39 bodies with combined certificates, typical durations fell into 6 to 12, 13 to 18, 19 to 24 and above 24 months with shares of 8, 51, 31 and 10 percent (Austrian National Public Health Institute et al., 2026b). Over all distributions of typical durations within these brackets, the standard deviation across bodies lies between 2.57 and 8.23 months if the open bracket ends at 36 months, and between 2.57 and 11.12 months if it ends at 48 months; the lower bound places every body at the point of its bracket closest to the common mean, and the upper bounds come from a grid search over masses at the bracket edges. Bodies differ in size, so these unweighted figures describe the dispersion of body-level practice and not the distribution across files.

D.9 Summary

No alternative parameterisation, distributional choice or modelling simplification reverses the three main findings (Table D.9). The positive variance gap is the most stable result; the sequencing comparison depends on how much a completed file shortens and narrows the European clock and on financing frictions.

E EXTENDED DISCUSSION

E.1 The valley of death and clearance-time variance

Research on the valley of death analyses financing gaps mainly through information asymmetries and the structure of capital markets (Murphy and Edwards, 2003; Parrish et al., 2019). Clearance-time variance adds a dimension, since the commercialisation date is itself widely distributed at the moment of investment. An investor who understands a device's clinical promise still cannot price the project precisely when the clearance distribution is wide or shifting, and the difference between a known interval of 13 to 18 months and an interval of 13 to 24 months or longer can move a project

Table D.9 Sensitivity of the main results to modelling choices

Variation	Effect on main results
Distribution family	CV estimates shift by up to about 13%; ordering of EU and US variance unchanged
Exponential service	Sojourn times overstated by about 44–59%; general service retained
Sequencing	US-first never optimal with an unchanged European file; with a completed file preferred where the achievable saving exceeds the requirement (Table B.1)
Interval D of 3–5 years	Differentials and sequencing comparison invariant; levels shift
Effective sample size $N/2$ to $2N$	Interval widths change; point estimates and rankings unchanged
CV band 0.25–0.45	$P(\tau > 24)$ for Class IIb rises from 5.0 to 13.8%; shortfall beyond 18 months from 1.04 to 2.24 months

across a financing threshold, especially for capital-constrained SMEs. Regulators can mitigate uncertainty in this phase (Bonnin Roca and O'Sullivan, 2022), and on the present account reducing σ_{EU} , as distinct from the mean delay, is the relevant form of that mitigation. Complexity of products has been associated with the difficulty of crossing the valley of death (Dean et al., 2022); complex devices also require more extensive dossiers, face more refusals and populate the upper tail of the processing-time distribution. Internal regulatory preparation can reduce refusal-driven tail events, but it cannot remove the queue dynamics that generate the variance gap, which is why the policy response has to be institutional.

Financing arithmetic makes the problem concrete. A developer whose funds last for the median interval between financing rounds, 22.9 months, faces a probability of between 10 and 41 percent that European certification outlasts them, as far as the survey brackets can tell, against 0.7 percent for a 510(k) clearance and 13.8 percent for an original premarket approval in the file-level records; expected shortfall beyond that horizon lies between 0.11 and 1.66 months for European certification if the open bracket ends at 36 months, and equals 0.04 and 2.41 months on the two US routes (Table 2 of the main text). Investors can price that exposure through larger financing rounds, staged commitments or a preference for markets with shorter tails, and each of these responses weighs most on firms without internal funds.

Evidence that manufacturers have moved launch priority away from the EU (MedTech Europe, 2025; Svempe, 2024) admits competing readings that the design cannot exclude as joint contributors. Pandemic disruption of European manufacturing SMEs overlaps the MDR transition (Juergensen et al., 2020); geopolitical risk pricing affects cross-Atlantic capital flows; long-run shifts in market volume reduce the weight of the EU for some device categories; and heterogeneity in discount rates is not captured by an aggregate δ . The variance channel is therefore one of several non-exclusive explanations, and separating their weights requires matched firm-level records of patent filings, FDA submissions and MDR applications. What the evidence supports is compatibility with the mechanism, not identification.

E.2 Consequences for the patent system

When effective EU terms are short and uncertain, firms have reason to redirect development towards products with faster US clearance, towards protection mechanisms other than patents, or towards drug–device combinations in which the medicinal component can obtain supplementary protection. A compositional shift of this kind reduces both the number of standalone devices developed for the EU first and the disclosure through which patents diffuse knowledge; configurational research on intellectual-property strategies describes the corresponding choice among protection mechanisms (Ayerbe et al., 2024). In the language of innovation systems, the notified-body bottleneck is a governance problem in which a constraint at one node propagates into the valuation of complementary actors (Granstrand and Holgersson, 2020; Granstrand, 2024). The effects reach beyond the EU perimeter as well, through global markets and supply chains that depend on European certification (Marley et al., 2025), and through the attractiveness of European research environments for device development (Zubiaur et al., 2025).

E.3 Restoring term lost to clearance: legal position and design parameters

Standalone devices are excluded from supplementary protection certificates for textual reasons. Regulation (EC) No 469/2009 covers products protected by a patent and authorised as medicinal products under the medicinal-products legislation (European Parliament and Council, 2009), and the Court of Justice has held that CE certification under the device regime cannot be treated as such an authorisation, even for a device incorporating a medicinal substance (Court of Justice of the European Union, 2018). The certificate’s duration equals the period between the patent filing and the first marketing authorisation in the Union, reduced by five years and capped at five years (Article 13(1) and (2)), and applications are filed with the competent industrial-property office of the member state that granted the patent (Article 9(1)). Its economic rationale, restoring term consumed by mandatory clearance, applies to standalone devices in the same form, which is why the question of a device-specific instrument recurs.

Welfare calibration and legal design of such an instrument are developed elsewhere (Peters, 2026, forthcoming). The present analysis adds one design parameter. A mean-based formula restores the average loss, whereas the variance channel implies that the loss is concentrated in the upper tail of the clearance distribution. An extension that grows with the excess of a device’s clearance time over a class-specific reference percentile, published from the notified-body monitoring data that Implementing Regulation (EU) 2026/977 introduces (European Commission, 2026), would respond to dispersion as well as to mean delay. The consumer side has to enter the same calculation, since additional exclusivity delays price competition, and device markets differ from pharmaceutical markets in life-cycle length and in the availability of substitutes; any extension should be calibrated to that device-specific welfare arithmetic.

E.4 Geographic and product-category incidence

Firm size is one of several dimensions along which incidence varies. Mechanically, the variance channel implies heavier incidence where access to notified bodies runs through few designated bodies and where clinical-investigation infrastructure is weaker. Conditional certification, a procedural instrument that could shorten access for high-risk devices addressing unmet needs, has seen very limited use (Dobrzynska et al., 2025), and the EU lacks a full counterpart to the accelerated US pathways for breakthrough devices (Gupte et al., 2025). Devices with complex dossiers and thin regulatory precedent therefore populate the upper tail, so the variance gap falls hardest on novel product categories. Policy that corrects only the firm-size dimension, for example through fee reductions, leaves the geographic and categorical incidence in place; a distributional audit along all three dimensions becomes feasible once record-level data are available.

E.5 Boundary conditions

Clearance-time variance is decision-relevant in a defined set of settings. Market entry must require mandatory approval that cannot be bypassed. Approval must consume residual protection term in a way that is not fully restored afterwards; automatic and complete restoration would reduce the mechanism to a problem of mean delay. Variance must be material relative to residual term and discount rate, since a five-month mean with four months of dispersion, the 510(k) case, differs structurally from a mean of eighteen months whose file-level dispersion no public record measures. Firms must also face risk-adjusted financing or opportunity costs, so that variance in timing becomes variance in project value instead of averaging out in a diversified portfolio. In-vitro diagnostics under the IVDR share the architecture, drug–device combinations escape it only for their medicinal component, and environmental permitting or aviation type certification display the same features of stochastic approval clocks, finite residual protection and risk-adjusted valuation. Whether the mechanism operates there is an empirical question for which the medical-device case supplies a methodological template.

Fig. E.1 sets out the architecture that the argument attributes the dispersion to.

Fig. E.2 states the same comparison as a decision tree for use at the project-valuation stage.

F REPRODUCTION PACKAGE

Every quantity in the article that rests on computation rather than on quotation of a source is regenerated by a single script, `fit_cv.py`, from data files that accompany it. The package contains one row per FDA decision with its fiscal year, submission type, third-party indicator and duration in days, separately for the 15,706 substantially equivalent premarket notifications, the 173 granted De Novo requests and the 311 original premarket approvals; the same premarket notification records with the per-applicant submission count, the strictly backward-looking count of prior decisions, the running submission order, the applicant country and the advisory committee, under an anonymous applicant identifier and without applicant names; the seventy parsed Federal Register notices with the exemption, submission and approval dates they state




Fig. E.1 The MDR conformity-assessment system as an architecture that distributes time risk. Solid arrows: governance and submission flows. Dashed arrows: dispersion of clearance times passed from the capacity-constrained population of notified bodies to manufacturers. Schematic, following the definition of innovation systems in Granstrand and Holgersson (2020)




Fig. E.2 Sequencing decision tree derived from Proposition 1 of the main text. $V(F)$ is the value of an EU submission with clearance distribution F , net of the submission cost and of the expected cost of waiting beyond the runway

and the patent number where one is named; the 161 device notices with the patents they name; the granted extensions from the patent office list; the notified-body survey series; the bracket definitions of Tables A.1 and A.2; the receipt-cohort records of Table D.1, one row per substantially equivalent decision received in fiscal years 2018 to 2022, with its dates, controls and applicant key; the applicant crosswalk on which that key rests; the Eurostat figures on enterprises and net turnover by size class; and a table of every externally sourced parameter with its value, the source and the page or figure it comes from. Values quoted from surveys and from data providers are listed in that table and not recomputed.

Running `python3 fit_cv.py` in the package directory writes to standard output and creates no files. It prints, in order, the throughput diagnostic of the notified-body record; the

bracket fits of Tables A.2 and A.3 for all three families, with the validation against the file-level 510(k) data; the moments, quantiles, tail probabilities and expected shortfalls of the four US populations and of the European fit; the phase medians of the Federal Register notices; the cost ratios by size class; the required reductions of mean European clearance time for every combination of US pathway, cost of capital, cost basis and refinancing cost; the firm-size fits, the mixture decomposition and the crossing point; the queueing quantities; the granted extensions and the comparison of notice and database dates; the least-squares, quantile and within-applicant regressions of log clearance time on applicant experience, with standard errors clustered by applicant throughout; and the six receipt-cohort estimates of Table D.1 under both year specifications.

A copy of the package is deposited at figshare under <https://doi.org/10.6084/m9.figshare.33683083> and released under a CC BY 4.0 licence. The archive there has the SHA-256 digest `a9a853495b620b3ee67e3a16d8eb6888f40f17678a193ae0f4fe7377e63a497c`, and a manifest inside it covers its seventeen other files, so a reader can establish before and after extraction that the copy being run is byte-for-byte the one described here. Running `sha256sum -c SHA256SUMS.txt` in the unpacked directory checks the seventeen; running `python3 fit_cv.py` and comparing the output with `expected_output.txt` checks the results. Three things about the package should be stated plainly. The script does not repeat the retrieval step: the FDA extracts, the Federal Register notices and the extension list are supplied as files with their access dates, and the code recomputes from them rather than downloading them again. The receipt-cohort records derive from the openFDA device 510(k) export archived on 12 September 2026, whose SHA-256 is recorded in the manifest and whose FDA data status is 7 September 2026; the file supplied here is the subset the estimation uses, not the whole export. And the applicant key is a reconstruction rather than an FDA field. Firm identity is not published: the export carries the applicant string, in which one firm appears under several spellings and legal forms. The crosswalk assigns 23,389 raw names appearing in decisions from 1996 onwards to firm identifiers, reproduces the identifier, the backward-looking prior count and the running submission number of all 15,702 analysis cases without exception, and marks every identifier created for a name outside that sample as newly generated. Reproducing the identifiers is a necessary condition and not a proof that the rule is the only one consistent with the record, and the crosswalk carries no claim about corporate control, renaming or acquisition.

G EFFECTIVE TERM AT RISK: PROOFS, BOUNDS AND COMPUTATIONS

G.1 Definitions and coherence

Let $q_u(\tau)$ denote the quantile function of clearance time. Expected shortfall $ES_\alpha = (1 - \alpha)^{-1} \int_\alpha^1 q_u du$ is monotone, translation-equivariant, positively homogeneous and subadditive, and therefore a coherent risk measure in the sense of Artzner et al. (1999), whereas the quantile itself is not subadditive in general (Acerbi and Tasche, 2002). Subadditivity matters in the present setting because devices are pooled into

Table G.1 Tail concealment under pooling, review phase of technical documentation assessments

α	q_α^L	q_α	q_α^S	ES_α^L	ES_α	ES_α^S	$q_\alpha^S - q_\alpha$	$ES_\alpha^S - ES_\alpha$
0.75	15.85	17.92	23.40	19.67	27.12	36.04	5.48	8.93
0.90	19.45	25.00	34.34	23.05	36.61	48.11	9.34	11.50
0.95	21.98	32.03	43.20	25.52	45.22	58.02	11.16	12.80

Note: Months; q_α and ES_α refer to the pooled distribution with $p = 0.6$; log-normal fits to the firm-size shares of MedTech Europe (2024, Fig. 8); pooled quantiles by root finding and pooled expected shortfall by numerical integration of the pooled quantile function. The last two columns give the months by which pooled statistics understate SME review time in the tail, equal to the protection they overstate.

portfolios and survey statistics pool across bodies and firm types. Effective term at risk and expected term shortfall follow through the affine map $T = T^{\max} - D - \tau$, which reverses the direction of risk and preserves the orderings below. Upper-tail dominance, the condition of Proposition 3 of the main text, is a local form of first-order stochastic dominance (Shaked and Shanthikumar, 2007).

G.2 Proof of Proposition 3

Assume continuous distributions with positive densities; with atoms the argument carries over to generalised inverses. Let $\alpha \geq \alpha_0$. Because $F_S(q_\alpha^S) = \alpha \geq F_S(x_0)$, the SME quantile satisfies $q_\alpha^S \geq x_0$, so upper-tail dominance gives $F_L(q_\alpha^S) \geq F_S(q_\alpha^S) = \alpha$ and therefore $F(q_\alpha^S) = pF_L(q_\alpha^S) + (1-p)F_S(q_\alpha^S) \geq \alpha$, which implies $q_\alpha \leq q_\alpha^S$. In the same way $q_\alpha^L \geq x_0$, hence $F_S(q_\alpha^L) \leq F_L(q_\alpha^L) = \alpha$ and $F(q_\alpha^L) \leq \alpha$, which implies $q_\alpha \geq q_\alpha^L$. The ordering holds at every level $u \geq \alpha$, so integrating the three quantile functions over $[\alpha, 1]$ yields $ES_\alpha^L \leq ES_\alpha \leq ES_\alpha^S$. $\square$

At the fitted moments of the technical-documentation review phase ($\mu_L = 13.37$, $\sigma_L = 4.63$, $\mu_S = 18.66$, $\sigma_S = 13.06$ months, $p = 0.6$) the group distributions cross once, at $x_0 = 10.16$ months, where both equal 0.259; above that point SME review times dominate. The order is asserted for levels strictly above α_0 : at α_0 itself a distribution with an atom at the crossing point can reverse it, and the proof in Section G assumes finite first two moments in each group. The crossing is not a technicality here, because below it a larger share of SME files is already finished, which is what the reported shares show in the 4 to 8 month bracket. Table G.1 gives the ordered quantities above the crossing point.

G.3 Distribution-free bounds from bracket data

Let reported shares $\hat{p}_k$ fall on brackets $[a_k, b_k)$, $k = 1, \dots, K$, with cumulative shares P_k , and close an open last bracket at a cap C ; the residual statutory term is the natural maximum, since clearance after expiry leaves no protection. For α with $P_{k^*-1} < \alpha \leq P_{k^*}$, every distribution that reproduces the shares has $q_\alpha \in [a_{k^*}, b_{k^*}]$ and

$$ES_\alpha \geq \frac{(P_{k^*} - \alpha)a_{k^*} + \sum_{k > k^*} \hat{p}_k a_k}{1 - \alpha}, \quad (G.1)$$

$$ES_\alpha \leq \frac{(P_{k^*} - \alpha)b_{k^*} + \sum_{k > k^*} \hat{p}_k b_k}{1 - \alpha}.$$

The quantile function is non-decreasing and confined to $[a_k, b_k]$ on $(P_{k-1}, P_k]$, so point masses at the lower edges attain the lower bound and point masses approaching the upper edges approach the upper bound. No distributional assumption enters the bounds, which form the identified region in the

sense of Manski (2003). Table G.2 applies them to the European survey shares and sets the result against the US records, where no bounds are needed because every decision date is published. Even the median of European typical durations is identified only within a bracket, between 13 and 18 months, and expected shortfall at the ninetieth percentile is bounded below by 24 months while its upper bound equals whatever cap is assumed for the open bracket. On the US side the same quantities are known to the month, and that contrast is what the table is meant to show.

G.4 Tail-indexed restoration and the value of reporting boundaries

An extension equal to the delay beyond a threshold R has expected cost $E[(\tau - R)^+]$ per dossier; with $R = q_\alpha$ it equals $(1 - \alpha)(ES_\alpha - q_\alpha)$. In the US records that cost is known exactly, at 0.49 months of additional protection per 510(k) dossier and 1.93 months per original premarket approval for $\alpha = 0.9$, concentrated on the one dossier in ten that needs it. For the European shares the same expression gives 0.32 months at the fitted distribution, with distribution-free bounds of -0.00 and 1.20 months at a cap of 36 months, since $(x - R)^+$ is non-decreasing and bracket data bound the cost between $\sum_k \hat{p}_k(a_k - R)^+$ and $\sum_k \hat{p}_k(b_k - R)^+$. The width of the identified region for ES_α equals $(1 - \alpha)^{-1}[(P_{k^*} - \alpha)(b_{k^*} - a_{k^*}) + \sum_{k > k^*} \hat{p}_k(b_k - a_k)]$; for the European shares at $\alpha = 0.9$ it is 12.0 months with a cap at 36 months and 36.0 months with a cap at 60 months, against zero in the US records. Reporting design, in other words, decides whether the quantity that matters for financing risk can be observed at all.

H SUPPLEMENTARY TABLES AND FIGURES REFERENCED IN THE MAIN TEXT

Tables H.1, H.2, H.3, H.4 and H.5 and Figs. H.1, H.2, H.3, H.4 and H.5 carry material that the main text refers to but does not need in full: the notified-body survey record on which the throughput diagnostic rests, what the current reporting template identifies about the tail, the effective-term-at-risk measures with their distribution-free bounds, the robustness summary, the glossary, and five figures, namely the throughput diagnostic, the causal chain of the argument, the level and spread of clearance time by pathway, the mixture-variance decomposition of the review phase, and the stylised contrast between a predictable and a dispersed clearance regime with which the article opens.

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)

Table G.2 Quantile and expected-shortfall bounds for the European shares, against exact US values (months)


Record	α	q_α fit	q_α bounds	ES_α fit	Lower	$C=36$	$C=60$
EU MDR, survey shares	0.50	17.61	[13, 18]	21.77	18.92	25.32	30.12
EU MDR, survey shares	0.75	20.85	[19, 24]	24.41	21.00	28.80	38.40
EU MDR, survey shares	0.90	24.28	[19, 24]	27.49	24.00	36.00	60.00
FDA 510(k)	0.50	4.37	exact	8.61		exact	
FDA 510(k)	0.90	10.38	exact	15.22		exact	
FDA 510(k)	0.95	14.03	exact	18.38		exact	
FDA PMA originals	0.50	10.48	exact	22.37		exact	
FDA PMA originals	0.90	27.04	exact	45.81		exact	
FDA PMA originals	0.95	37.36	exact	61.00		exact	

Note: European rows use the shares of the twentieth notified-body survey and the log-normal fit of Table A.2; C is the cap assumed for the open bracket above 24 months. US rows are computed from the file-level records and require no bounds. Reproduced by the script of Section F.

Table H.1 Notified-body survey record for the MDR: cumulative applications and certificates (sum of annexes)



Reference date	Round	Applications	Certificates	Ratio
Oct 2022	EC data	8,120	1,990	0.25
Mar 2023	S1	10,818	2,951	0.27
Jun 2023	S4	13,177	3,899	0.30
Oct 2023	S6	17,846	5,599	0.31
Feb 2024	S8	20,424	6,978	0.34
Jun 2024	S10	25,987	8,900	0.34
Oct 2024	S12	28,037	10,497	0.37
Feb 2025	S14	28,489	12,177	0.43
Jun 2025	S16	32,920	14,831	0.45
Oct 2025	S18	33,107	17,507	0.53
Feb 2026	S20	32,898	18,010	0.55

Note: Medium dataset of the survey, applications and certificates summed across conformity-assessment annexes, as charted in the comparison published with round 20 (Austrian National Public Health Institute et al., 2026b); “EC data” denotes the series collected by the European Commission before the study rounds began. The October 2025 row carries the values revised in that comparison; the eighteenth round had first published 33,175 applications and 17,549 certificates (Austrian National Public Health Institute et al., 2026a). The decline of 209 applications in February 2026 reflects changes in counting methods at a few bodies, flagged in the source, and for three bodies the application count derives from the bimonthly small dataset, which reports a different aggregate (31,902 applications lodged by 28 February 2026). Fifty-two bodies were designated under the MDR at that date.

**Fig. H.1** Throughput diagnostic from the notified-body survey record (Table H.1). Circles: cumulative ratio of certificates to applications. Squares: marginal ratio of new certificates to new applications for the four intervals inside the October 2022 to February 2024 window (shaded); later marginal ratios are dominated by revisions of earlier counts and are not interpreted

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>

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*. Com-

**Fig. H.2** The variance channel from regulatory architecture to distributional incidence. Arrows correspond to the formal mappings of Section 4 of the main text**Fig. H.3** Clearance time by pathway. Markers show means and bars the 10th to 90th percentiles. US values come from FDA file-level records; the EU MDR row for bodies uses the log-normal fit to the twentieth survey, the row for firms the average technical-documentation certification time reported by MedTech Europe (2024), and the MDD row the range implied by the statement that MDR certification takes twice as long. The dotted line marks the median interval between financing rounds of 696 days

missioned 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

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>

Bañas, L., Dawid, H., Randrianasolo, T. A., Storn, J., & Wen, X. (2022). Numerical approximation of a system of Hamilton–

Table H.2 What duration reporting identifies about the tail

Reporting requirement	Identified	Not identified
Median duration, as in Art. 4(2) of Implementing Regulation (EU) 2026/977	Centre of the distribution	Upper tail, effective term at risk, expected shortfall, runway probabilities
Median and 90th percentile	Effective term at risk at the 90 percent level	Expected shortfall and probabilities beyond the 90th percentile
Shares in brackets with an open top interval, as in the current surveys	Quantiles within one bracket; a lower bound for expected shortfall	Upper bound for expected shortfall, which depends on the assumed cap
Shares in brackets closed by an upper boundary B	Expected shortfall within bounds whose width shrinks as B falls	Distribution within brackets

Note: Based on the bounds of Section G. The compliance share of Art. 4(2) refers to the maximum timelines of individual assessment activities and does not identify the tail of total duration. With the pooled MDR shares and a boundary at 48 months that contains all mass above 30 months, the width of the identified region for expected shortfall at the 90 percent level falls from 115.2 to 14.4 months.

Table H.3 Tail quantiles and expected shortfall of clearance time: identified values and distribution-free bounds

Record	α	q_α (months)		ES_α (months)		
		Value	Bounds	Value	Lower	Upper, cap 36 / 60
EU MDR, survey, log-normal fit	0.50	17.61	[13, 18]	21.77	18.92	25.32 / 30.12
EU MDR, survey, log-normal fit	0.90	24.28	above 24, open	27.49	24.00	36.00 / 60.00
FDA 510(k), file level	0.90	10.38	exact	15.25	exact	exact
FDA 510(k), file level	0.95	14.03	exact	18.40	exact	exact
FDA PMA originals, file level	0.90	27.04	exact	46.36	exact	exact
FDA PMA originals, file level	0.95	37.36	exact	58.69	exact	exact

Note: Months. Quantiles and expected shortfall use the inverse empirical distribution function for the US records. No effective term in years is reported, because that would require the date from which the term-bearing patent runs, which is the earliest application in its chain of benefit and is not established here. EU bounds hold for every distribution of typical durations per body that reproduces the survey counts; at $\alpha = 0.90$ the cumulative share at 24 months is 0.897, so the quantile falls in the open upper bracket and has no upper bound in the data.

- Jacobi–Bellman equations arising in innovation dynamics. *Journal of Scientific Computing*, 92, 54. <https://doi.org/10.1007/s10915-022-01892-x>
- 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>
- 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>
- 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.
- 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>
- Dean, T., Zhang, H., & Xiao, Y. (2022). The role of complexity in the Valley of Death and radical innovation performance. *Technovation*, 109, 102160. <https://doi.org/10.1016/j.technovation.2020.102160>
- 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>

Table H.4 Robustness of the main quantities to documented alternatives

Quantity	Central case	Documented alternative	Effect
Mean EU certification	18.17 months, survey fit	21.8 months, firm-reported TDA (MedTech Europe, 2024)	Gap to the 510(k) route rises from 1.05 to 1.36 years
510(k) records	All substantially equivalent decisions, 5.51 months, CV 0.80	Traditional submissions without third-party review, 6.22 months, CV 0.70	Gap to the EU narrows by 0.7 months; exceedance of 22.9 months 0.8 instead of 0.7 percent
PMA period	FY2016–2025, 14.72 months	FY2021–2025, 16.24 months, CV 1.09	Gap to the EU narrows by 1.5 months
Filing to submission D	Median 8.89 years	Quartiles 6.90 and 13.39 years	Residual term of 13.10 or 6.61 instead of 11.11 years; differences between pathways unchanged
Cost of capital δ	9.57 percent	7.54 and 10.57 percent	Required saving behind a 510(k) review changes by at most 0.9 months
Cost of a further round c	0 to 0.90 years of revenue	Upper and lower values as shown	Required saving behind a 510(k) review 5.3 to 5.5 months at $c = 0$ and 3.7 to 3.9 months at $c = 0.90$ with fee-based costs
EU dispersion	CV 0.26 across bodies	CV 0.25 and 0.45 at the same mean	Probability of more than 24 months 10.5 and 19.4 percent

Note: Months unless stated otherwise. Required savings for all combinations of cost of capital, cost basis and firm size in Section B.5; distributional families in Section D. Reproduced by the script of Section F.

Table H.5 Glossary of terms and symbols

Term	Meaning in this article
Appropriability	Extent to which an innovator can capture the returns from an invention, here through the patent term that remains after market authorisation.
Between-group component	Term $p(1-p)(\mu_L - \mu_S)^2$ of pooled variance that arises only because firm types differ in mean clearance time (Proposition 2 of the main text).
Clearance time τ_j	Random time from regulatory submission to certification or clearance in jurisdiction j , with mean μ_j and standard deviation σ_j .
Clearance-time variance H_j	Variance of effective patent term, equal to σ_j^2 as long as patent expiry does not truncate the distribution (Definition 2 of the main text).
Coefficient of variation (CV)	Standard deviation divided by the mean; a scale-free measure of dispersion estimated from reported duration brackets.
Conformity assessment	Procedure under the MDR by which a notified body verifies that a device meets the regulatory requirements before it may be placed on the market.
Effective patent term T_j	Statutory term of twenty years minus the interval from filing to submission and minus clearance time (Definition 1 of the main text).
Effective term at risk ETaR $_{\alpha}$	Protection retained at the α -quantile of clearance time (equation (8) of the main text).
Expected term shortfall ETS $_{\alpha}$	Protection retained on average by the slowest share $1 - \alpha$ of dossiers, based on expected shortfall (equation (8) of the main text).
Expected shortfall $E[(\tau - R)^+]$	Expected number of months by which clearance exceeds the financing interval R (equation (3) of the main text).
Applicant experience class	Number of substantially equivalent premarket notifications an applicant obtained in fiscal years 2021–2025, in four classes; the regressor of equation (4) of the main text.
Financing interval R	Time for which the funds raised in one round last, measured by the median interval between funding rounds; $P(\tau > R)$ is the probability that clearance outlasts it.
Cost of a further round c	Cost of an additional financing round in years of revenue, measured as median dilution times the ratio of enterprise value to sales (equation (5) of the main text).
Implementing Regulation (EU) 2026/977	Commission act of 4 May 2026 setting uniform procedural requirements, maximum timelines and monitoring duties for notified bodies.
MDD	Medical Devices Directive 93/42/EEC, the predecessor regime to the MDR.
MDR	Regulation (EU) 2017/745 on medical devices.
MDUFA	Medical Device User Fee Amendments, under which the FDA commits to review-time goals financed by user fees.
Mixture-variance identity	Law of total variance applied to two firm types; splits pooled variance into within-group and between-group terms (equation (7) of the main text).
Notified body	Conformity-assessment body designated by a member state to certify devices under the MDR.
Partial identification	Situation in which data bound a quantity without pinning it down; here, tail measures bounded by reported bracket shares (Table H.3).
PMA and 510(k)	US premarket approval for the highest-risk devices, and premarket notification based on substantial equivalence for moderate-risk devices.
Pollaczek–Khinchine relation	Queueing formula linking mean waiting time to traffic intensity and to the dispersion of service times (equation (2) of the main text).
Quantiles P10, P50, P90	Clearance times below which 10, 50 and 90 percent of dossiers fall under a reconstructed distribution (Table 2 of the main text).
Residual statutory term T^{res}	Patent term remaining at regulatory submission, $T^{\text{max}} - D$; 11.11 years at the median interval of the Federal Register sample.
Sequencing condition	Comparison of the discounted value of a European submission after completing the file during the US review with the value of submitting the file as it stands (equation (6) of the main text).
Supplementary protection certificate (SPC)	EU instrument extending protection for medicinal products whose authorisation consumed patent term; unavailable for standalone devices.
Throughput ratio $\theta(t)$	Cumulative certificates divided by cumulative applications at a survey reference date.
Traffic intensity ρ	Ratio of the arrival rate of dossiers to the service capacity of the assessing bodies.

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.

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>

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>

Gupte, A., Nitave, R., & Gobburu, J. (2025). Regulatory landscape of accelerated approval pathways for medical devices in




Fig. H.4 Mixture-variance decomposition of the review phase of MDR technical documentation assessments (months and squared months), from the firm-size distributions in MedTech Europe (2024, Figure 8) with log-normal bracket fits and the SME share of certificates in MedTech Europe (2022). The within-group term dominates because the SME distribution is far wider




Fig. H.5 Why variance matters independently of average delay. Two stylised clearance regimes share a mean of 15 months (dotted line) but differ in spread (log-normal, CV 0.22 and 0.45). The dispersed regime sends a larger share of devices into the long-wait tail beyond 24 months (shaded), where more patent term has elapsed before market entry

the United States and the European Union. *Frontiers in Medical Technology*, 7, 1586070. <https://doi.org/10.3389/fmedt.2025.1586070>

Huusko, 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>

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>

Kleinrock, L. (1975). *Queueing systems, Vol. I: Theory*. New York: Wiley.

Manski, C. F. (2003). *Partial identification of probability distributions*. New York: Springer. <https://doi.org/10.1007/b97478>

Marley, T., Barrow, D., & McDermott, O. (2025). Global market and supply chain implications of the European Union medical device regulation 2017/745 as amended by 2023/607. *BMC Health Services Research*, 26, 155. <https://doi.org/10.1186/s12913-025-13907-y>

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 (2022). *Survey report: availability of medical devices in connection to MDR implementation*. Brussels: MedTech Europe.

MedTech Europe (2025). *IVDR & MDR survey results 2024: highlights report*. Brussels: MedTech Europe, January 2025.

Murphy, L. M., & Edwards, P. L. (2003). *Bridging the valley of death: transitioning from public to private sector financing*. NREL/MP-720-34036. Golden, CO: National Renewable Energy Laboratory.

Parrish, M. C., Tan, Y. J., Grimes, K. V., & Mochly-Rosen, D. (2019). Surviving in the valley of death: opportunities and challenges in translating academic drug discoveries. *Annual Review of Pharmacology and Toxicology*, 59, 405–421. <https://doi.org/10.1146/annurev-pharmtox-010818-021625>

Peters, A. (2026). 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. (forthcoming). A device protection certificate for medical devices under EU law: doctrinal feasibility and legislative design after Boston Scientific. *GRUR International*.

Shaked, M., & Shanthikumar, J. G. (2007). *Stochastic orders*. New York: Springer. <https://doi.org/10.1007/978-0-387-34675-5>

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>

Zubiaur, J., Rumoroso Cuevas, J. R., Rodríguez Mateos, M. C., et al. (2025). Medical device research update after the adoption of the EU legislation (MDR). *REC: Interventional Cardiology*, 7(4), 255–263. <https://doi.org/10.24875/RECICE.M25000529>