

Medicine as a Safety-Critical System

A Probabilistic Risk Assessment of 50 U.S. Diseases

Lev Neymotin, PhD
N5Digital, Inc.

February 3, 2026

Abstract

Medicine routinely makes decisions that involve irreversible harm, high societal cost, and significant uncertainty. Yet, unlike other safety-critical domains—such as nuclear power, aerospace engineering, or chemical processing—medicine lacks a unified, explicit framework for probabilistic risk assessment (PRA). Instead, medical decision-making relies on a patchwork of outcome metrics, cost-effectiveness analyses, clinical risk scores, and expert judgment, none of which provide a transparent, population-level accounting of preventable risk.

This position paper proposes the application of Probabilistic Risk Assessment to medicine as a complementary analytical framework for prioritizing research, guiding regulatory evaluation, and identifying opportunities for systemic risk reduction. We introduce a simple, transparent risk metric defined as the product of disease frequency (expressed as U.S. cases per capita per year) and graded consequence severity, capturing not only mortality but also chronic disability, functional impairment, recurrence, and societal burden.

Using this framework, we analyze fifty common diseases in the United States, spanning acute, chronic, and procedural conditions. The resulting risk-ordered landscape reveals structural patterns that are often obscured in conventional medical analyses: high population risk frequently arises from conditions with low individual severity but extreme frequency, while a distinct subset of diseases exhibits substantial operator-dependent outcome variance, making them particularly amenable to risk reduction through precision-enforcing interventions.

We further introduce a formal method for estimating achievable risk reduction (ΔRisk) through the suppression of operator variance, drawing direct analogies to safety-credit methodologies used in engineering PRA. While the numerical values presented are illustrative rather than definitive, the framework itself is intended to be general, extensible, and data-agnostic.

This work does not seek to replace existing clinical or epidemiological tools. Rather, it argues that medicine, as a safety-critical system, would benefit from adopting an explicit PRA perspective—one that makes risk tradeoffs visible, comparable, and structurally accountable. A focused case study on Dupuytren's contracture is presented in the Appendix to demonstrate how this framework can inform precision-driven procedural redesign.

1. Introduction

1.1 Medicine as a Safety-Critical System

Modern medicine operates at a scale and complexity that rivals other safety-critical systems. Each year, hundreds of millions of medical interventions are performed in the United States alone, many of them invasive, irreversible, and dependent on real-time human execution. These interventions carry not only the possibility of therapeutic benefit but also the risk of death, permanent disability, chronic pain, recurrence, and long-term societal cost.

Despite these characteristics, medicine is rarely analyzed using the same formal risk frameworks that govern other domains where failure is unacceptable. In fields such as nuclear power generation or aerospace engineering, the presence of low-probability, high-consequence events has led to the widespread adoption of Probabilistic Risk Assessment (PRA): a structured methodology that explicitly quantifies risk as the product of event frequency and consequence severity, and that guides investment in safety features, procedural design, and regulatory oversight.

Medicine, by contrast, has developed a rich set of tools for measuring outcomes, costs, and individual patient risk, but lacks a unifying framework for evaluating **population-level, preventable risk** in a transparent and comparable manner. As a result, decisions about research priorities, procedural standardization, and technology adoption are often made implicitly, guided by historical precedent, professional norms, or economic incentives rather than by explicit risk accounting.

1.2 Existing Medical Risk Metrics and Their Limits

It is important to emphasize that medicine is not devoid of quantitative analysis. Measures such as Quality-Adjusted Life Years (QALYs), Disability-Adjusted Life Years (DALYs), cost-effectiveness ratios, and disease-specific risk scores play critical roles in clinical research and health policy. However, these tools are designed to answer fundamentally different questions. [1,2]

Outcome-based metrics quantify the impact of disease after it has occurred. Cost-effectiveness analyses optimize economic efficiency within predefined intervention sets. Clinical risk scores estimate prognosis for individual patients. None of these approaches are intended to function as a system-level probabilistic risk framework analogous to PRA in engineering, where the central objective is to identify, compare, and reduce avoidable risk across an entire system.

Most notably, existing medical frameworks rarely isolate **operator-dependent variance** as a first-class contributor to risk. In many medical procedures, outcomes depend not only on disease biology but also on the precision, consistency, and execution skill of the practitioner. In safety-critical engineering, such variability is treated as a design flaw to be engineered out of the system wherever feasible. In medicine, it is often accepted as an intrinsic feature of practice.

1.3 Purpose and Scope of This Paper

This paper argues that medicine would benefit from adopting a Probabilistic Risk Assessment perspective—not as a replacement for existing tools, but as a complementary framework for decision-making at the system level. Specifically, we propose a simple, transparent risk metric defined as:

$$\text{Risk} = \text{Frequency} \times \text{Consequence Severity}$$

where frequency is expressed as U.S. cases per capita per year, and consequence severity is graded on a standardized scale that incorporates mortality, disability, chronic burden, recurrence, and societal impact.

Using this framework, we examine fifty common diseases in the United States, selected to span a broad range of frequencies, severities, and intervention types. The objective is not to produce definitive rankings, but to demonstrate how an explicit PRA-style approach reveals structural patterns in medical risk that are otherwise difficult to see.

In addition, we introduce a formal method for estimating potential risk reduction through **precision enforcement**—the use of engineering, imaging, and decision-support systems to reduce operator-driven outcome variance. This approach mirrors established PRA methodologies used to justify safety investments in other high-risk industries.

A detailed case study of Dupuytren’s contracture is provided in the Appendix to illustrate how this framework can be applied to a specific procedural disease and connected to concrete proposals for precision-guided intervention. [5]

It should be noted that in specific procedural domains, including Dupuytren’s contracture, practitioners hold a range of views regarding the necessity and value of additional precision-enforcing tools beyond established clinical techniques. These views reflect accumulated experience, practical constraints, and differing interpretations of procedural risk. The present work does not seek to adjudicate these positions. Rather, it aims to provide a probabilistic framework within which such differences in judgment can be evaluated systematically, using explicit assumptions about frequency, consequence severity, and reducibility of operator-dependent variance.

1.4 Positioning and Limitations

This work is intentionally framed as a **position paper**. The numerical values presented are illustrative and rely on simplified assumptions intended to preserve transparency and comparability rather than epidemiological precision. The goal is to establish a coherent analytical structure that can later be refined with higher-resolution data, demographic stratification, and disease-specific modeling.

The central claim is not that medicine has ignored risk, but that it has done so implicitly. By introducing an explicit PRA framework, this paper seeks to make medical risk tradeoffs visible, discussable, and subject to systematic improvement—bringing medicine conceptually closer to other domains where safety is treated as an engineering problem rather than an article of faith.

2. Probabilistic Risk Assessment: From Engineering to Medicine

2.1 Origins and Purpose of Probabilistic Risk Assessment

Probabilistic Risk Assessment (PRA) emerged in engineering disciplines where system failure carries the potential for irreversible, large-scale harm. Its development was driven not by theoretical elegance, but by necessity: as engineered systems grew more complex, deterministic safety margins proved insufficient to account for low-probability, high-consequence events. [5,10]

In its canonical form, PRA evaluates risk as the product of two components: the frequency (or probability) of an initiating event and the severity of its consequences. This formulation allows heterogeneous risks—rare catastrophes and frequent minor failures alike—to be expressed on a common scale. PRA does not seek to eliminate uncertainty; rather, it structures uncertainty in a way that enables rational comparison, prioritization, and intervention.

In practice, PRA serves several tightly coupled functions in safety-critical fields. First, it makes implicit risk tradeoffs explicit by expressing heterogeneous hazards in a common quantitative language. Second, it enables prioritization by identifying where limited resources can yield the greatest reduction in total system risk. Finally, PRA provides direct design guidance by highlighting vulnerabilities that arise from excessive reliance on human perfection, thereby motivating engineered safeguards that suppress outcome variance.

Importantly, PRA is not a predictive tool in the narrow sense. It does not forecast specific accidents. Instead, it characterizes a risk landscape and guides decisions about system design, operational constraints, and investment in safety features.

2.2 PRA in Safety-Critical Engineering Domains

PRA is now a foundational element of safety governance in multiple industries, including nuclear power, aerospace, chemical processing, and large-scale infrastructure systems. In these contexts, PRA informs regulatory requirements, licensing decisions, and the allocation of research and development resources.

A defining feature of engineering PRA is its treatment of human action. Human operators are recognized as necessary but imperfect components of complex systems. Consequently, modern safety design aims not to demand flawless execution, but to **reduce the consequences of error and suppress outcome variance** through engineered constraints, automation, redundancy, and real-time feedback.

This philosophy reflects a core insight: when system outcomes depend critically on individual operator skill, the system itself is inadequately designed. PRA provides a quantitative basis for identifying such vulnerabilities and justifying investments that shift safety from human-dependent to system-enforced mechanisms.

2.3 Structural Parallels Between Engineering Systems and Medical Care

Medicine shares many defining characteristics with other safety-critical systems. Medical interventions are frequently irreversible, time-sensitive, and executed under uncertainty. Adverse outcomes may involve permanent injury, loss of function, or death. Moreover, the aggregate societal impact of medical harm is substantial, both economically and ethically.

Despite these parallels, medicine has historically been treated as categorically distinct from engineered systems. Clinical decision-making is framed primarily as an exercise in individualized judgment, informed by training, experience, and evolving evidence. While this emphasis on professional expertise is essential, it has also obscured the extent to which medical outcomes depend on system design choices rather than on biology alone.

In particular, many medical procedures exhibit high **operator-dependent variance**. Outcomes may vary significantly across practitioners performing nominally identical interventions, even when treating similar patients. In engineering PRA, such variance would be interpreted as a signal that the system relies excessively on human precision and insufficiently on structural safeguards.

The absence of a PRA framework in medicine has allowed this variance to persist largely unquantified. As a result, it is difficult to compare the risk contribution of different diseases, procedures, or intervention strategies on a common basis, or to justify investments aimed specifically at reducing execution-dependent risk.

2.4 Why PRA Has Not Historically Been Applied to Medicine

The limited adoption of PRA-style thinking in medicine is not due to a lack of mathematical or conceptual compatibility. Rather, it reflects historical, cultural, and institutional factors.

First, medicine evolved as a craft long before it became a large-scale, systematized enterprise. The centrality of individual clinical judgment remains a defining feature of professional identity. Second, ethical concerns about quantifying harm—particularly non-fatal harm—have made explicit risk tradeoffs

uncomfortable, even when they are implicitly made in practice. Third, existing outcome-based metrics have often been treated as sufficient proxies for risk, despite addressing different analytical questions.

Finally, the complexity and heterogeneity of medical data have discouraged the use of simplified models. However, engineering PRA routinely operates under similar constraints, using conservative assumptions and bounding analyses to inform decisions in the presence of uncertainty. Precision is refined over time; structure comes first.

2.5 Rationale for a Medical PRA Framework

Applying PRA to medicine does not imply reducing clinical care to formulas or diminishing the role of professional judgment. Instead, it introduces a complementary perspective focused on **system-level risk** rather than individual outcomes.

A medical PRA framework enables several capabilities that remain underdeveloped in current practice. It allows heterogeneous diseases to be compared on a common risk scale, clarifying whether population burden is driven primarily by frequency or by severity. It also makes it possible to isolate procedural domains in which outcome variance is disproportionately influenced by execution quality rather than disease biology. Together, these insights provide a transparent basis for evaluating and justifying investment in precision-enforcing technologies aimed at reducing preventable harm.

In this sense, PRA serves as a bridge between clinical practice and system design. It allows medicine to ask not only whether an intervention works, but whether the **structure of care delivery itself is optimized to minimize preventable harm.**

The sections that follow introduce a simplified PRA-style risk metric tailored to medicine and apply it to a representative set of common diseases in the United States. The intent is not to provide definitive risk estimates, but to demonstrate how an explicit probabilistic framework reveals structural features of medical risk that are otherwise difficult to detect.

3. A Probabilistic Risk Metric for Medicine

3.1 Conceptual Requirements for a Medical PRA Metric

In engineering PRA, the purpose of a risk metric is not predictive precision but **comparability and prioritization**. A useful metric must allow heterogeneous failure modes—rare catastrophes and frequent minor failures—to be evaluated on a common scale, even when data are incomplete or uncertain [4,5].

For medicine, an analogous risk metric must satisfy several conditions simultaneously. It must be population-based rather than patient-specific, explicitly combining how often harm occurs with how severe that harm is. At the same time, it must remain simple enough to preserve transparency while being sufficiently structured to support comparison across disease categories. Finally, it must accommodate both fatal and non-fatal outcomes, including chronic disability, functional impairment, and recurrence.

These requirements closely mirror those that motivated the original development of PRA in nuclear and chemical engineering [5,10].

3.2 Definition of Risk Components

In this study, medical risk is defined as:

$$R = F \times CS$$

where:

- **Frequency (F)** is the number of disease cases per capita per year in the United States.
- **Consequence Severity (CS)** is a dimensionless graded score (1–100) reflecting the magnitude of harm associated with a typical case.

This formulation follows the foundational PRA principle that **risk is the product of likelihood and consequence**, a relationship that has remained stable across engineering domains for decades [4,5].

3.3 Frequency: Cases per Capita per Year

Frequency is expressed as **U.S. cases per capita per year** to normalize diseases with vastly different incidence scales.

This normalization allows episodic conditions, such as influenza or the common cold, to be treated consistently alongside chronic diseases represented by clinically significant annual events. It also permits procedural diseases to be incorporated without privileging mortality-based definitions, ensuring that non-fatal but function-limiting conditions contribute appropriately to population-level risk.

Engineering PRA routinely relies on frequency estimates that are approximate, conservative, and periodically revised as data improve [4,10]. The same philosophy is adopted here: frequency estimates are intended to be **order-of-magnitude accurate** and comparable, not exact.

3.4 Consequence Severity Scale

Consequence Severity (CS) is graded on a **1–100 scale**, grouped into four conceptual bands:

Band 1: 76–100 (Severe): death, profound disability, permanent loss of independence, or catastrophic societal cost.

Band 2: 51–75 (Medium): chronic disease with lasting functional impairment, sustained medical dependence, or major productivity loss.

Band 3: 26–50 (Mild): temporary but clinically significant morbidity requiring medical care.

Band 4: 1–25 (Low): inconvenience or minor functional limitation with minimal long-term impact.

This approach deliberately extends beyond mortality. Engineering PRA similarly incorporates non-lethal consequences such as environmental damage, economic loss, and long-term system disruption [5,6]. In medicine, chronic pain, loss of function, recurrence, and dependency are ethically and economically significant harms and must be included in any population-level risk assessment.

3.5 Interpretation and Limits of the Metric

The risk metric used here does **not** represent individual patient risk, prognosis, or expected utility. Instead, it reflects **aggregate population burden**, making visible the cumulative impact of diseases that may appear benign at the individual level but dominate total risk due to frequency.

As with early PRA applications in engineering, this first-pass abstraction is intentionally simple. Its value lies not in numerical precision but in the **structural clarity it brings to decision-making** [4,5].

4. Dataset and Methodology

4.1 Selection of Diseases

Fifty diseases were selected to provide a representative cross-section of medical conditions encountered in the United States.

Diseases were selected based on high prevalence or incidence, clinical relevance across age groups, and inclusion of acute, chronic, and procedural conditions. Additional consideration was given to relevance for public health expenditure and workforce productivity. The intent was not exhaustive coverage, but structural demonstration across a representative cross-section of medical risk.

The intent is not exhaustive coverage but **structural demonstration**—analogous to representative system modeling in engineering PRA [5].

4.2 Case Definition and Normalization

Diseases were categorized according to their dominant clinical pattern. Acute episodic diseases, such as influenza or gastroenteritis, were counted as episodes per year. Chronic diseases, including diabetes and osteoarthritis, were represented by clinically significant annual cases or new diagnoses. Procedural diseases were represented by treated cases or procedural episodes. This mixed approach reflects practical constraints in medical data reporting and mirrors the integration of heterogeneous event definitions common in engineering PRA.

This mixed approach reflects practical constraints in medical data reporting. Engineering PRA similarly integrates heterogeneous event definitions under a common risk framework [4,10].

4.3 Data Philosophy

The analysis prioritizes **consistency over granularity**. While medical epidemiology distinguishes carefully between incidence, prevalence, and lifetime risk, such distinctions—while essential for clinical research—can obscure population-level comparisons when applied inconsistently.

Here, the objective is to construct a **coherent risk landscape**, not to resolve disease-specific epidemiological debates. This is consistent with PRA practice in non-medical domains, where early models often precede high-resolution data refinement [5,6].

4.4 Assumptions and Constraints

Several simplifying assumptions are made explicitly. Demographic stratification by age, sex, or socioeconomic status is not included, comorbidities are not modeled, and consequence severity is averaged across disease presentations. Interactions between diseases are also not considered. These limitations are acknowledged as areas for future refinement, but they do not invalidate the comparative purpose of the analysis.

These limitations are acknowledged as areas for future extension. However, they do not invalidate the comparative purpose of the analysis, just as early PRA models remained valuable despite similar simplifications [10].

5. Results: Risk Ranking of 50 Common U.S. Diseases

Applying the probabilistic risk framework defined above, fifty common diseases in the United States were analyzed and ranked according to population-level risk, defined as the product of disease frequency and consequence severity.

The objective of this analysis is not to establish definitive numerical rankings, but to reveal the structural distribution of medical risk across disease categories. To emphasize these structural features, selected results are summarized in Table 5.1, which highlights how different classes of disease contribute to total population risk through distinct combinations of frequency, severity, and procedural dependence. The complete set of disease-specific estimates is provided in Appendix B.

Table 5.1 – Risk Structure by Disease Class (Illustrative)

Disease Class	Typical Frequency	Typical Severity	Dominant Risk Driver
Acute infectious illnesses	Very high	Low–Medium	Frequency
Chronic degenerative diseases	Medium	Medium–High	Cumulative burden
Cardiovascular events	Medium	High	Severity
Procedural diseases	Low–Medium	Medium	Operator-dependent variance

(Full disease-specific risk estimates are provided in Appendix B.)

5.1 Overview of the Risk-Ordered Landscape

The resulting risk-ordered landscape differs markedly from conventional disease hierarchies based on mortality or cost alone.

Two dominant patterns emerge:

1. **High-frequency, low-severity diseases account for a substantial fraction of total population risk.**
2. **Several moderate-frequency, moderate-severity conditions rival or exceed traditionally “serious” diseases in aggregate risk contribution.**

This mirrors a well-known phenomenon in engineering PRA, where frequent minor failures often dominate total system risk, while rare catastrophes dominate public attention [5,10].

5.2 Frequency-Dominated Risk

Conditions such as the common cold, gastroenteritis, sinus infections, and influenza rank highly due to extreme frequency despite relatively low consequence severity. Individually, these conditions are rarely life-threatening; collectively, they account for enormous societal burden through lost productivity, healthcare utilization, and cumulative morbidity.

This finding is not surprising mathematically, but it is often underappreciated in medical policy discussions that focus disproportionately on mortality endpoints.

5.3 Severity-Dominated Risk

At the opposite end of the spectrum, diseases such as cancer, ischemic heart disease, and stroke maintain high risk rankings primarily due to severity rather than frequency. These conditions remain critical targets for prevention and treatment, but the PRA framing clarifies **why** they matter—and how they compare quantitatively to other sources of harm.

5.4 Emergence of Procedural Risk Clusters

A notable result of the ranking is the clustering of diseases where outcomes depend heavily on **procedural execution** rather than disease biology alone. These conditions—including orthopedic, cardiovascular, ophthalmologic, and interventional diseases—exhibit moderate-to-high risk combined with substantial operator-dependent variance.

In engineering PRA, such clusters would immediately prompt design review and consideration of variance-suppressing interventions [4,5]. Their identification here motivates the focused analysis developed in the subsequent sections.

5.5 Implications of the Ranking

The purpose of the ranking is not to declare certain diseases “more important” than others in an ethical sense. Rather, it provides a **quantitative basis for asking different questions**:

- Which risks are driven primarily by frequency?
- Which by severity?
- Which by procedural execution?
- Where is engineering intervention likely to yield the greatest reduction in preventable harm?

6. Procedural Diseases and Operator-Driven Risk

6.1 Defining Procedural Diseases in a PRA Context

Within a probabilistic risk framework, a *procedural disease* is defined as a condition for which a substantial portion of patient outcome variance arises not from disease biology alone, but from the **execution quality of a medical intervention**. This includes surgical procedures, catheter-based interventions, injections, implant placement, and other techniques where correct targeting, force, timing, or spatial precision materially influence outcomes.

In engineering PRA, systems whose performance depends critically on human execution are recognized as intrinsically high-risk unless compensatory design measures are implemented [5,10]. Operator-dependent variance is treated not as an unavoidable reality, but as a **risk contributor** that can—and should—be quantified and reduced where feasible.

Applying this logic to medicine reframes procedural skill not as a purely individual attribute, but as a **system property** subject to design optimization.

6.2 Distinguishing Biological Risk from Procedural Risk

Not all medical risk is procedural. Many diseases progress according to biological mechanisms largely independent of clinician action. However, in procedural diseases, outcomes often depend on a layered sequence:

1. **Biological baseline** (disease state, patient variability),
2. **Decision accuracy** (indication selection, timing),
3. **Execution precision** (technical performance of the procedure),
4. **Collateral effects** (off-target damage, complications),
5. **Recurrence or re-intervention** driven by incomplete or imprecise treatment.

Only the first layer is irreducibly biological. The remaining layers are, in principle, **system-controllable**. In traditional clinical evaluation, these layers are often conflated. PRA explicitly separates them, enabling identification of the fraction of total risk attributable to operator execution rather than disease inevitability [4,5].

6.3 Procedural Diseases Identified in the 50-Disease Dataset

Applying the above definition to the dataset analyzed in Section 5 reveals a distinct subset of conditions where operator-driven variance plausibly contributes a significant share of total risk. These include:

- **Surgical-dominant conditions**
(e.g., cataract disease, colorectal cancer, breast and prostate cancer resections), where margin control, tissue handling, and complication avoidance affect long-term outcomes.
- **Interventional cardiovascular conditions**
(e.g., ischemic heart disease, stroke interventions), where time-sensitive execution, device positioning, and access technique influence morbidity and mortality.
- **Orthopedic and musculoskeletal conditions**
(e.g., osteoarthritis requiring joint replacement, selected low back pain interventions), where alignment, fixation, and procedural judgment affect function and revision rates.
- **Injection- and targeting-dependent conditions**
where small spatial errors can produce disproportionate harm or reduced efficacy.

In these diseases, PRA flags a structural issue: **the same nominal intervention can produce widely different outcomes across operators**, even when clinical indications are similar.

6.4 Operator Variance as a First-Class Risk Contributor

In engineering PRA, variability introduced by human action is not treated as random noise. It is explicitly modeled, bounded, and mitigated through system design—using physical constraints, redundancy, automation, and feedback mechanisms [5,10].

Medicine, by contrast, has traditionally addressed operator variance primarily through training, credentialing, and experience accumulation. While necessary, these approaches do not eliminate variance; they merely attempt to shift its distribution.

From a PRA perspective, reliance on operator perfection represents an **unstable safety strategy**. When adverse outcomes are strongly execution-dependent, the system remains vulnerable even when average performance improves. PRA therefore emphasizes reducing the *sensitivity* of outcomes to execution error rather than attempting to eliminate error entirely [4,6].

6.5 Why Procedural Diseases Are PRA-Relevant

Procedural diseases occupy a unique position in the medical risk landscape:

- Their **frequency** is often moderate rather than extreme.
- Their **consequence severity** is non-trivial but not uniformly catastrophic.
- Their **risk is disproportionately reducible** through system-level intervention.

This combination makes them especially important from a PRA standpoint. In nuclear safety, analogous conditions—moderate-frequency initiating events with controllable consequences—are prioritized for design improvement because they offer the greatest return on safety investment [10].

The same logic applies in medicine. Identifying procedural diseases allows risk-reduction efforts to move beyond incremental biological advances toward **engineering-driven precision**, where improvements in execution reliability can yield outsized reductions in aggregate harm.

6.6 Implications for Risk Reduction Strategy

Recognizing procedural diseases as a distinct class changes the nature of the risk-reduction question. Instead of asking solely whether a treatment is effective, PRA encourages asking:

- How much of the observed harm is execution-dependent?
- How sensitive are outcomes to small procedural deviations?
- Can variance be suppressed through design rather than training alone?

In several procedural fields, including those involving injection- or targeting-dependent interventions, practitioners may reasonably conclude that existing techniques already achieve acceptable outcomes within current standards of care. From a PRA perspective, such conclusions represent implicit risk assessments based on experience rather than explicit quantification. The framework developed here is intended to complement, not replace, that experience by making underlying assumptions about acceptable risk and variance reduction transparent and comparable across conditions.

These questions lead naturally to the concept of **precision enforcement**, developed in the next section, where operator-driven risk is treated as a quantifiable and reducible component of total medical risk.

7. Risk Reduction Through Precision Enforcement

7.1 Concept of Precision Enforcement

In safety-critical engineering, risk reduction is rarely pursued by demanding higher levels of human perfection. Instead, system designers seek to **reduce the sensitivity of outcomes to human error** by introducing constraints, guidance, automation, and feedback mechanisms. This approach—variously implemented through physical interlocks, redundancy, sensor fusion, and control logic—has proven more reliable than training-based mitigation alone [5,10].

In this paper, the term **precision enforcement** is used to describe an analogous strategy in medicine: the deliberate use of system-level design to constrain procedural execution, guide operator action, and suppress outcome variance attributable to human technique. Precision enforcement does not eliminate clinical judgment, but it limits the degrees of freedom within which judgment is exercised, thereby reducing preventable harm.

Examples of precision enforcement may include, but are not limited to:

- imaging-guided targeting systems,
- mechanical alignment or depth-limiting devices,
- real-time feedback and confirmation mechanisms,
- standardized execution pathways supported by decision systems.

From a PRA perspective, the defining feature of precision enforcement is not the specific technology employed, but its **effect on the risk equation**: reducing the contribution of operator-dependent variance to total risk.

7.2 Decomposing Medical Risk in Procedural Diseases

To quantify the impact of precision enforcement, total medical risk must be decomposed into components that are biologically intrinsic and components that are procedurally contingent. Let total risk for a disease be:

$$R = F \times CS$$

as defined in Section 3.

For procedural diseases, total risk can be conceptualized as comprising two additive contributions:

1. **Biological risk**, driven by disease pathology and patient variability.
2. **Procedural risk**, driven by execution quality, targeting accuracy, collateral damage, and recurrence related to incomplete or imprecise intervention.

Only the second component is directly addressable through precision enforcement. PRA requires that this component be isolated and quantified, even if only approximately, to assess the potential benefit of system redesign [4,5].

7.3 Δ Risk: Formal Definition

We define Δ Risk as the achievable reduction in total risk resulting from the suppression of operator-driven variance through precision enforcement.

Let:

- F = cases per capita per year,
- CS = consequence severity,
- q = fraction of cases that involve a procedural intervention where execution quality materially affects outcome,
- k = fraction of consequence severity attributable to operator-dependent variance within those cases,
- r = fraction of that operator-dependent variance that can realistically be reduced through precision enforcement.

Then the reducible portion of total risk is:

$$\Delta R = F \times CS \times q \times k \times r$$

and the residual risk after intervention becomes:

$$R' = R - \Delta R = F \times CS \times (1 - qkr)$$

This formulation is directly analogous to **risk reduction credit** calculations used in engineering PRA to justify safety features, procedural redesigns, or automation [5,10].

7.4 Interpretation of Parameters

Each parameter in the Δ Risk expression has a clear conceptual meaning:

- **Procedural exposure fraction (q)**
Represents how much of the disease burden is actually subject to procedural execution. Many diseases are mixed; only a subset of cases may involve interventions where precision matters.
- **Operator-dependent severity fraction (k)**
Captures how much harm arises from execution variability rather than unavoidable disease progression. Complications, collateral damage, incomplete treatment, and recurrence are included here.
- **Achievable reduction factor (r)**
Represents the extent to which system design can realistically suppress execution-driven

variance. PRA practice emphasizes conservative estimation of such factors, especially in early analyses [4,6].

Crucially, PRA does not require these parameters to be known with high precision at the outset. Their purpose is to structure uncertainty, not to eliminate it [5].

7.5 Scenario-Based Risk Reduction Analysis

In engineering PRA, early-stage evaluations often use **scenario bands**—conservative, moderate, and aggressive—to bound plausible outcomes and guide decision-making under uncertainty [10]. The same approach is appropriate here.

For example:

- A **conservative scenario** might assume that only a small fraction of cases are procedural, that operator variance contributes modestly to harm, and that precision enforcement yields limited improvement.
- A **moderate scenario** reflects partial procedural prevalence and meaningful but incomplete variance suppression.
- An **aggressive scenario** applies where procedures are dominant drivers of outcome and precision enforcement directly addresses the primary failure modes.

The objective is not to assert a single correct value, but to determine whether ΔRisk is **non-trivial across plausible assumptions**. If ΔRisk remains meaningful even under conservative assumptions, the case for system-level intervention is strengthened.

7.6 PRA Implications for Medical Design and Policy

Introducing ΔRisk shifts the medical risk conversation in several important ways:

1. **From outcome comparison to risk reducibility**
Diseases are evaluated not only by how harmful they are, but by how much harm is preventable through design.
2. **From training optimization to system optimization**
Operator excellence remains important, but PRA emphasizes reducing reliance on perfect execution.
3. **From implicit tradeoffs to explicit justification**
Investment in precision technologies can be evaluated using the same logic that governs safety investments in other high-risk industries.

It is important to emphasize that identification of a non-zero ΔRisk does not imply that existing practice is unsafe or deficient. In engineering PRA, many safety enhancements are implemented in systems already considered acceptable, based on the recognition that modest reductions in variance can yield meaningful aggregate benefits. Applied to medicine, this perspective suggests that precision-enforcing approaches may be evaluated as incremental risk-reduction options rather than as corrections of flawed practice.

In engineering contexts, failure to pursue available risk reduction once identified through PRA carries regulatory and ethical implications [4,10]. While medicine operates under different constraints, the underlying logic—that preventable risk should be systematically addressed—remains applicable.

7.7 Transition to Case Study

The framework developed in this section is intentionally general. Its purpose is to enable disease-agnostic reasoning about procedural risk and its reducibility. In the Appendix, this framework is applied to a specific procedural disease—Dupuytren’s contracture—to demonstrate how PRA concepts translate into concrete design and intervention proposals.

8. Implications for Medical Research, Regulation, and Ethics

8.1 Implications for Medical Research and Development

Applying a probabilistic risk assessment framework to medicine alters how research priorities can be evaluated. Traditional medical R&D is often guided by biological novelty, unmet clinical need as defined by severity, or market considerations. While each of these factors is important, none explicitly incorporate **population-level preventable risk** as a decision criterion.

A PRA-based perspective introduces an additional dimension: **risk reducibility**. Diseases characterized by moderate severity but high frequency, or by substantial operator-dependent variance, may represent more effective targets for system-level intervention than conditions defined solely by catastrophic outcomes. In engineering PRA, such conditions frequently justify investment because incremental improvements yield disproportionate reductions in total system risk [5,10].

For medical research, this suggests a complementary prioritization strategy in which investment decisions consider not only therapeutic potential, but also the magnitude of risk that can realistically be reduced through improved system design, procedural standardization, or precision-enforcing technologies.

8.2 Implications for Regulatory Evaluation

Regulatory frameworks in medicine primarily evaluate safety and efficacy at the level of individual products or procedures. While post-market surveillance captures adverse events, there is limited formal mechanism for assessing whether **available system-level risk reductions** are being pursued once a procedure becomes standard practice.

In safety-critical engineering, PRA is routinely used to justify regulatory requirements for design modifications when preventable risk exceeds accepted thresholds [4,10]. Although medicine operates under different ethical and practical constraints, the underlying principle—that avoidable harm should be systematically addressed—remains applicable.

A medical PRA framework could inform regulatory evaluation by:

- highlighting procedural domains with disproportionate execution-dependent risk,
- providing a structured basis for evaluating precision-guidance technologies,
- supporting risk-informed updates to standards of care.

Such an approach would not mandate specific technologies, but would encourage **explicit justification** when known sources of preventable risk remain unaddressed.

8.3 Implications for Clinical Practice and Training

This framework does not diminish the importance of clinical expertise or professional judgment. Rather, it reframes their role within a broader system context. PRA distinguishes between risks that are unavoidable and risks that persist due to system design choices.

In procedural diseases, continued reliance on training alone as the primary mitigation strategy may be insufficient. Engineering experience demonstrates that training-based risk reduction is inherently unstable when outcomes remain sensitive to small execution errors [5]. PRA instead encourages reducing the dependence of outcomes on individual performance by narrowing allowable execution variability.

For clinicians, this perspective supports the adoption of tools and workflows that enhance consistency and reliability, not as constraints on autonomy, but as safeguards against preventable harm.

8.4 Ethical Considerations

Explicit risk quantification raises legitimate ethical concerns, particularly regarding the valuation of non-fatal harm and the comparison of disparate disease burdens. However, these comparisons are already made implicitly through funding decisions, reimbursement policies, and technology adoption.

A PRA-based framework does not introduce ethical tradeoffs; it **makes them visible**. By articulating assumptions and uncertainties explicitly, PRA supports more transparent and accountable decision-making.

From an ethical standpoint, tolerating preventable risk without explicit justification is increasingly difficult to defend, particularly when system-level interventions could reduce harm without compromising care quality. In this sense, PRA aligns with core ethical principles of non-maleficence and justice.

8.5 Limitations and Cautions

Several cautions are warranted. PRA should not be interpreted as a deterministic decision rule or a substitute for clinical judgment. Numerical rankings must be understood as approximations, sensitive to assumptions about frequency and severity.

Moreover, not all operator-dependent variance is readily reducible. Some aspects of medical care resist standardization due to patient heterogeneity or evolving disease dynamics. PRA is most effective when applied selectively, focusing on domains where execution precision is both critical and controllable.

These limitations underscore the role of PRA as a **decision-support framework**, not a prescriptive algorithm.

8.6 Broader Implications for Healthcare Systems

At a system level, introducing PRA into medicine encourages a shift from retrospective harm analysis to **prospective risk management**. Instead of asking why adverse events occurred after the fact, PRA emphasizes identifying where harm is most likely and most preventable before it occurs.

This shift has implications for how healthcare systems evaluate new technologies, allocate resources, and define quality improvement goals. Over time, a PRA-informed approach could contribute to a more coherent alignment between clinical practice, engineering design, and public accountability.

9. Discussion

9.1 Reframing Medical Risk Without Replacing Medicine

The central contribution of this paper is not a new epidemiological dataset or a competing clinical metric, but a **reframing of how medical risk is conceptualized at the system level**. By importing Probabilistic Risk Assessment (PRA) into medicine, the analysis introduces a transparent structure for combining disease frequency and consequence severity into a single population-level risk measure.

This reframing does not invalidate existing medical tools. Clinical trials, observational studies, QALYs, DALYs, and disease-specific risk scores remain essential for understanding efficacy, prognosis, and cost. PRA operates at a different level: it asks how preventable harm is distributed across the healthcare system and where intervention is most likely to reduce aggregate risk. In this sense, PRA is complementary rather than competitive.

Importantly, PRA does not claim to resolve uncertainty. Engineering PRA was developed precisely because uncertainty is unavoidable in complex systems [5,10]. Its value lies in structuring uncertainty so that decisions can be justified, revisited, and improved over time.

9.2 Comparison With Existing Medical Risk and Burden Frameworks

It is instructive to contrast the PRA approach with commonly used medical burden metrics. QALYs and DALYs integrate morbidity and mortality into unified measures of health loss, but they are primarily designed for **economic evaluation and resource allocation**, often after interventions are already defined. They do not explicitly address the question of *where system redesign could most effectively reduce risk*.

Similarly, clinical risk scores are patient-level tools optimized for prediction and triage. They provide little guidance for population-level investment in procedural precision or system reliability. Failure Mode and Effects Analysis (FMEA) and Root Cause Analysis (RCA) address process safety locally, but they are episodic, retrospective, and not easily comparable across diseases or institutions [1,2].

PRA differs in scope and intent. It is explicitly **prospective**, comparative, and system-oriented. It asks not only how bad outcomes are, but how often they occur and how reducible they are through design. This distinction explains why PRA has become foundational in other safety-critical domains and suggests why its absence in medicine represents a conceptual gap rather than a settled methodological choice.

9.3 Operator Variance as a System Design Issue

One of the most consequential insights emerging from the PRA framing is the treatment of operator-dependent variance as a **system property**, not an individual failing. In engineering, systems that depend critically on flawless human execution are considered poorly designed, regardless of operator skill [5].

Medicine has traditionally managed operator variance through training, credentialing, and experience. While these measures improve average performance, they do not eliminate variance, nor do they protect against fatigue, rare errors, or the long tail of adverse outcomes. PRA makes this residual risk visible and provides a rationale for shifting from training-dominant mitigation to **design-dominant mitigation** where feasible.

This perspective does not diminish professional expertise. Rather, it acknowledges that even highly skilled practitioners operate within systems that can amplify or suppress error. PRA reframes the ethical

question from “Who made a mistake?” to “Why did the system allow this outcome to depend so strongly on individual execution?”

9.4 Interpreting the 50-Disease Risk Landscape

The analysis of fifty common U.S. diseases illustrates several structural features of medical risk. First, it confirms that high population risk often arises from conditions that are individually mild but extremely frequent. Second, it reveals that several diseases with moderate severity contribute risk comparable to traditionally “serious” diseases due to their prevalence.

Most importantly for this paper’s argument, the analysis highlights a subset of **procedural diseases** where outcome variance is plausibly reducible. These conditions do not dominate risk because they are biologically catastrophic, but because they combine frequency, non-trivial consequences, and execution sensitivity. In engineering PRA, such conditions are prime candidates for targeted risk reduction [10].

The specific rankings should not be interpreted as definitive. They are sensitive to assumptions about case definitions and severity scoring. However, the qualitative patterns—frequency dominance, severity dominance, and procedural clusters—are robust to reasonable variations in input parameters.

9.5 Limitations and Future Refinement

Several limitations of the present framework warrant discussion. The use of simplified frequency estimates and averaged severity scores necessarily obscures heterogeneity across patient populations. Demographic stratification, comorbidities, and disease interactions are not modeled. These omissions limit the framework’s use for fine-grained policy decisions.

Nevertheless, these limitations are consistent with early-stage PRA applications in engineering, where coarse models often precede data refinement and system redesign [5,10]. The intent here is to establish a **conceptual scaffold** upon which more detailed analyses can be built.

Future extensions could incorporate age- or sex-specific risk profiles, longitudinal accumulation of risk, and integration with real-world outcome registries. Importantly, such refinements would strengthen rather than alter the underlying PRA logic.

9.6 Anticipated Objections

A common objection to applying PRA in medicine is that healthcare is “too complex” or “too individualized” for such abstraction. Similar objections were raised in the early adoption of PRA in nuclear safety, where systems were also heterogeneous, context-dependent, and partially understood. PRA succeeded not by eliminating complexity, but by managing it explicitly.

Another concern is that quantitative risk ranking may oversimplify ethical judgments. However, PRA does not dictate ethical priorities; it exposes the tradeoffs that already exist. Ethical decision-making is weakened, not strengthened, when such tradeoffs remain implicit.

Finally, some may argue that existing safety and quality frameworks are sufficient. This paper does not dispute their value. It argues instead that they operate without a unifying probabilistic structure, limiting their ability to guide system-level design choices across diseases and institutions.

9.7 Positioning of This Work

This paper should be understood as a **position statement** advocating for the adoption of PRA principles in medicine. It does not propose regulatory mandates or prescribe specific technologies. Its aim is to introduce a language and logic that allow preventable medical risk to be discussed with the same rigor routinely applied in other safety-critical fields.

In doing so, it invites further empirical work, debate, and refinement. If medicine is to continue evolving as a large-scale, technology-mediated system, the explicit treatment of risk as a system property may become not only useful, but necessary.

10. Conclusion

This paper has argued that medicine, when viewed at the scale and complexity at which it now operates, meets the defining criteria of a safety-critical system. Like other such systems, it routinely balances uncertainty, irreversible harm, and societal cost. Yet unlike nuclear power, aerospace engineering, or chemical processing, medicine lacks a unified probabilistic framework for making these tradeoffs explicit and comparable.

By introducing a Probabilistic Risk Assessment (PRA) perspective to medicine, this work provides a structured method for combining disease frequency and consequence severity into a transparent population-level risk metric. Applied to fifty common diseases in the United States, this framework reveals structural features of medical risk that are not readily apparent in outcome-based or cost-based analyses alone. In particular, it highlights the disproportionate contribution of high-frequency conditions to total risk and identifies a subset of procedural diseases where outcome variance is plausibly reducible through system-level design.

The intent of this paper is not to supplant existing medical metrics or clinical judgment, but to complement them. PRA does not prescribe specific interventions or ethical priorities; rather, it makes implicit assumptions visible and subject to scrutiny. By isolating operator-dependent variance as a quantifiable risk contributor, the framework provides a rationale for shifting portions of medical risk management from training and experience toward precision-enforcing system design.

The numerical results presented here are illustrative and deliberately conservative. Their purpose is to demonstrate how a PRA structure can be applied to medicine, not to offer definitive rankings or policy mandates. As with early applications of PRA in engineering, refinement of data, modeling, and assumptions should be expected as the framework is tested, debated, and extended.

The broader implication is that medicine may benefit from adopting the same intellectual discipline that governs other safety-critical domains: treating preventable harm as a system property rather than an unavoidable byproduct of human imperfection. Doing so does not diminish the art of medicine; it strengthens its ethical foundation by aligning clinical excellence with system reliability.

In areas where practitioner judgment currently supports the sufficiency of existing techniques, PRA may still offer value by articulating why that judgment is reasonable and under what conditions alternative approaches might merit reconsideration.

The Appendices that follow provide a detailed case study of Dupuytren's contracture and the full disease risk tables used in this analysis. Together, they illustrate how probabilistic risk assessment can move

from abstract principle to practical insight, and how medicine might evolve toward a more explicit, accountable, and design-informed approach to safety.

Appendix A: Dupuytren's Contracture as a Procedural Risk Case Study

A.1 Rationale for Case Selection

Dupuytren's contracture provides a particularly clear illustration of the value of a Probabilistic Risk Assessment (PRA) framework applied to medicine. While the condition is rarely fatal, it is common, progressive, function-limiting, and frequently recurrent. Most importantly, treatment outcomes are **highly sensitive to procedural execution**, making it an archetypal *procedural disease* as defined in Section 6.

From a PRA perspective, Dupuytren's occupies a region of the medical risk landscape characterized by:

- moderate disease frequency,
- non-trivial functional consequences,
- substantial operator-dependent outcome variance,
- multiple available interventions with materially different precision requirements.

These features make Dupuytren's contracture an appropriate case study for examining **risk reducibility through precision enforcement**, rather than biological innovation alone.

A.2 Clinical Background (Brief)

Dupuytren's contracture is a fibroproliferative disorder of the palmar fascia leading to progressive digital flexion contractures. Disease progression is variable and may involve multiple cords and joints over time. Functional impairment increases nonlinearly as contracture severity advances, particularly when proximal interphalangeal joints are involved.

Current interventional treatments include:

- open or limited fasciectomy,
- percutaneous needle aponeurotomy,
- enzymatic cord disruption (e.g., collagenase injection).

All of these interventions involve **targeted mechanical or chemical disruption of pathological cords**, performed in close proximity to neurovascular structures.

A.3 Dupuytren's in the PRA Framework

A.3.1 Frequency (F)

For PRA purposes, Dupuytren's frequency is expressed as **treated cases per capita per year**, acknowledging that:

- not all diagnosed cases are treated immediately,
- recurrence often necessitates repeat interventions.

While precise epidemiological definitions vary, order-of-magnitude estimates place treated incidence on the order of **10^{-3} per person-year** in the U.S. population, placing Dupuytren's squarely in the moderate-frequency regime relevant for PRA screening.

A.3.2 Consequence Severity (CS)

Using the severity rubric defined in Section 3, Dupuytren's contracture is assigned a **CS in the mid-range (≈ 55)**, reflecting:

- loss of hand function,
- occupational and daily activity impairment,
- procedural complications,
- recurrence and cumulative morbidity.

While individual interventions are often labeled "minor," PRA emphasizes **aggregate and longitudinal burden**, particularly when recurrence is common.

A.3.3 Baseline Risk

Baseline population risk is therefore:

$$R_{Dup} = F \times CS$$

Although this value is lower than that of high-frequency acute illnesses, its **risk structure is qualitatively different**: a large fraction of harm arises from **how** the intervention is executed rather than from disease inevitability.

A.4 Operator-Dependent Risk Decomposition

Dupuytren's contracture exhibits several features that amplify operator-dependent variance:

1. **Targeting sensitivity**
Effective treatment requires accurate identification and engagement of pathological cords while avoiding adjacent structures.
2. **Collateral damage risk**
Off-target enzyme diffusion or mechanical injury may produce tendon rupture, nerve injury, or vascular compromise.
3. **Incomplete disruption and recurrence**
Imprecise targeting increases the likelihood of residual pathological tissue, accelerating recurrence and re-intervention cycles.
4. **Technique variability across operators**
Outcomes vary substantially with experience and technique, even within nominally standardized procedures.

From a PRA standpoint, these factors contribute to a high **operator-dependent severity fraction (k)** and a high **procedural exposure fraction (q)** relative to many other non-lethal conditions.

A.5 ΔRisk Analysis for Dupuytren's Contracture

Using the ΔRisk framework defined in Section 7:

$$\Delta R = F \times CS \times q \times k \times r$$

For Dupuytren's:

- q is high, as nearly all clinically significant harm occurs in the context of procedural intervention.
- k is substantial, reflecting complication, recurrence, and incomplete efficacy driven by execution.
- r is plausibly non-trivial, as targeting precision, force control, and confirmation of engagement are amenable to system-level design improvements.

Even conservative assumptions yield a meaningful ΔRisk , particularly when recurrence-driven re-intervention is incorporated by defining an **effective frequency**:

$$F_{eff} = F \times (1 + p_{re} \cdot m)$$

where p_{re} is the probability of recurrence and m the mean number of re-interventions.

In PRA terms, precision enforcement acts by **reducing both consequence severity per intervention and effective event frequency**, analogous to reducing initiating-event probability in engineered systems.

A.6 Link to the Dupuytren Precision Proposal

The author has proposed a precision-guided procedural system for Dupuytren's contracture intended to reduce targeting uncertainty and execution-dependent variance. This proposal, and the responses it has elicited from practicing clinicians, illustrate the range of views that currently exist regarding the necessity and feasibility of additional precision enforcement in this domain.

Regardless of the eventual disposition of any specific proposal, Dupuytren's contracture serves as a useful case study for demonstrating how probabilistic risk assessment can provide a neutral framework for examining procedural risk, practitioner judgment, and potential avenues for incremental improvement.

Appendix B: Risk Assessment of 50 Common U.S. Diseases

Table B1. Population-Level Probabilistic Risk Assessment of 50 Common U.S. Diseases
(Note: the risk values are order-of-magnitude PRA estimates.)

Disease	Disease Class	Frequency (cases per capita / year)	Consequence Severity (CS)	Risk ($F \times CS$)
Common cold	Acute	2.9	15	43.8
Acute gastroenteritis	Acute	0.44	35	15.3
Ischemic heart disease	Procedural	0.053	100	5.26
Seasonal influenza	Acute	0.117	45	5.26
Allergic rhinitis	Chronic	0.175	20	3.51
Acute sinusitis	Acute	0.102	25	2.56
Osteoarthritis	Procedural	0.041	60	2.45
Low back pain	Mixed / Procedural	0.047	50	2.34
Major depressive disorder	Chronic	0.029	75	2.19
COVID-19 (symptomatic)	Acute	0.029	70	2.05
Anxiety disorders	Chronic	0.044	40	1.75
GERD	Chronic	0.058	30	1.75
Acute pharyngitis	Acute	0.073	20	1.46

Disease	Disease Class	Frequency (cases per capita / year)	Consequence Severity (CS)	Risk (F × CS)
Dental caries (treated)	Procedural	0.088	15	1.32
Urinary tract infection	Acute	0.029	30	0.88
Acute bronchitis	Acute	0.029	25	0.73
Obesity (new)	Chronic	0.015	50	0.73
Periodontal disease	Procedural	0.029	25	0.73
Otitis media	Acute	0.029	20	0.58
Cancer (all types)	Procedural	0.0058	95	0.56
Skin cancer (treated)	Procedural	0.016	35	0.55
Cataract	Procedural	0.010	40	0.41
Hyperlipidemia	Chronic	0.0088	45	0.39
Type 2 diabetes (new)	Chronic	0.0044	85	0.37
Pneumonia	Acute	0.0044	80	0.35
Migraine (chronic)	Chronic	0.0073	40	0.29
Heart failure	Procedural	0.0029	90	0.26
Rheumatoid arthritis	Chronic	0.0038	70	0.27
Asthma	Chronic	0.0053	45	0.24
Anemia	Chronic	0.0058	40	0.23
Stroke	Procedural	0.0023	90	0.21
Hypertension (new)	Chronic	0.0029	60	0.18
Conjunctivitis	Acute	0.0175	10	0.18
Chronic kidney disease	Chronic	0.0020	85	0.17
COPD	Chronic	0.0018	80	0.14
Alzheimer's disease	Chronic	0.0015	95	0.14
Osteoporosis	Chronic	0.0020	65	0.13
Bipolar disorder	Chronic	0.0015	70	0.10
Peptic ulcer disease	Chronic	0.0015	55	0.08
Breast cancer	Procedural	0.00088	90	0.079
Prostate cancer	Procedural	0.00085	80	0.068
Lung cancer	Procedural	0.00070	95	0.067
Dupuytren's contracture	Procedural	0.00088	55	0.048
Epilepsy	Chronic	0.00058	70	0.041
Colorectal cancer	Procedural	0.00044	90	0.039
Gout	Chronic	0.00088	40	0.035
Endometriosis	Chronic	0.00058	55	0.032
Parkinson's disease	Chronic	0.00026	85	0.022
Multiple sclerosis	Chronic	0.000058	80	0.0047

References

1. **Wreathall J, Nemeth CP.** *Assessing risk: The role of probabilistic risk assessment (PRA) in patient safety improvement.* Qual Saf Health Care. 2004;13(3):206–212. Discusses PRA concepts and their potential relevance to healthcare as a complement to existing safety methods.
2. **Alemi F.** *Probabilistic risk analysis is practical.* Qual Manag Health Care. 2007;16(4):300–310. Shows how healthcare organizations can use PRA to analyze near-miss and sentinel events and prioritize risk reduction.
3. **Comden SC, et al.** *Using probabilistic risk assessment to model medication system risk.* In: *Risk Modeling and Patient Safety.* (Detailed ST-PRA application in long-term care settings.)
4. **U.S. Nuclear Regulatory Commission (NRC).** *Probabilistic Risk Assessment: PRA.* Official NRC description of PRA methodology and risk calculation fundamentals from safety-critical engineering (nuclear industry).
5. **Henley ER, Kumamoto H.** *Probabilistic Risk Assessment and Management for Engineers and Scientists.* Wiley-Interscience. Standard engineering textbook on PRA methodology and risk modeling.
6. **Scientific Advisory Board / EPA.** *Probabilistic Risk Assessment White Paper.* Environmental Protection Agency guidance on probabilistic methods and uncertainty characterization in risk assessment.
7. **PubMed Central Article:** *The role of probabilistic risk assessment (PRA) in patient safety.* Notes PRA's potential for analyzing adverse outcome causality and complements existing safety tools like FMEA/RCA.
8. **Probabilistic Risk Assessment (Wikipedia).** Summary of PRA definitions, key principles, and quantitative risk characterization (probability × consequence). Useful as an overview citation.
9. **ResearchGate overview:** *Probabilistic Risk Assessment (PRA).* Describes systematic quantitative risk analysis methodology and its use in engineering contexts.
10. **WASH-1400 (NUREG-75/014).** *The Reactor Safety Study.* Early foundational probabilistic risk assessment applied to nuclear reactor safety and development of PRA methods.