

Radiation Oncology Physics

Detailed High-Yield Reference | ABR Core Exam Prep

This reference expands each core physics concept with the governing formula, key variables, clinical application, a worked example, and a high-yield pitfall, organized by topic area following the general structure of the ABR physics core curriculum.

Atomic and Nuclear Physics

Atomic Structure and Characteristic X-Rays

Formula / Concept	Characteristic x-ray energy equals the difference in binding energy between the shell an electron falls from and the shell vacancy it fills.
Key Variables	Binding energy (shell-dependent, increases with atomic number Z); K, L, M shell notation.
Clinical Application	Characteristic x-rays are a minor component of the therapeutic megavoltage spectrum but are clinically important in kV imaging, where target material and filtration are chosen partly around characteristic emission lines.
Worked Example	A K-shell vacancy in a tungsten target is filled by an outer-shell electron, emitting a characteristic K-alpha x-ray around 59 keV, which appears as a discrete peak superimposed on the continuous bremsstrahlung spectrum used in kV imaging.
High-Yield Pitfall	Characteristic x-ray production requires the incident electron energy to exceed the binding energy of the specific inner-shell electron being ejected, not simply any arbitrary threshold energy.

Radioactive Decay Kinetics

Formula / Concept	$N = N_0 e^{-\lambda t}$, half-life $t_{1/2} = \ln 2 / \lambda$
Key Variables	N_0 = initial number of atoms; λ = decay constant; $t_{1/2}$ = half-life.
Clinical Application	Governs brachytherapy source strength decay over time; short half-life isotopes (such as Ir-192) require frequent source exchange, while long half-life isotopes (such as Cs-137) allow extended clinical use.
Worked Example	An Ir-192 source with a roughly 74-day half-life decays to about half its original activity after 74 days and about a quarter after 148 days, illustrating why HDR afterloader sources are typically exchanged on a schedule of a few months to keep treatment times clinically reasonable.
High-Yield Pitfall	Decay is an exponential, not linear, process — a common error is assuming activity decreases by equal absolute amounts per unit time rather than by a constant fraction each half-life.

Modes of Radioactive Decay

Formula / Concept	Alpha decay, beta-minus decay, beta-plus decay/electron capture, isomeric transition with gamma emission.
Key Variables	Each mode changes mass number and/or atomic number differently; gamma emission typically follows a preceding particulate decay that leaves the daughter nucleus in an excited state.
Clinical Application	Co-60 teletherapy relies on beta-minus decay to excited Ni-60, which then emits therapeutic gamma rays; brachytherapy isotopes decay by varying mechanisms that determine their emission spectrum and shielding requirements.
Worked Example	Co-60 decays by beta emission to excited Ni-60, which then emits two gamma rays at 1.17 and 1.33 MeV; the clinically quoted 'Co-60 energy' of 1.25 MeV is the average of these two distinct emissions.
High-Yield Pitfall	The commonly cited 1.25 MeV for Co-60 is an average of two separate gamma lines, not a single characteristic photon energy — a frequently tested point.

Photon Beam Production and Interactions

Bremsstrahlung X-Ray Production

Formula / Concept	Bremsstrahlung photons are produced when an incident electron decelerates in the electric field of a target nucleus, producing a continuous spectrum of photon energies up to the maximum electron energy.
Key Variables	Target atomic number Z ; incident electron kinetic energy.
Clinical Application	High- Z targets (such as tungsten) are used in linear accelerators and kV imaging tubes to maximize bremsstrahlung x-ray yield for a given electron energy.
Worked Example	Increasing the electron energy (kVp in diagnostic imaging, or nominal MV in therapy) increases both the quantity and the maximum energy of the bremsstrahlung photons produced, shifting the spectrum toward higher average energy.
High-Yield Pitfall	A clinical designation such as '6 MV' is a nominal beam-quality/accelerating-potential designation, not an exact measured maximum photon energy — the upper end of the bremsstrahlung spectrum is approximately 6 MeV, but the mean photon energy is substantially lower still, so never equate the MV label with the mean photon energy.

Linear Accelerator Beam Production

Formula / Concept	Electron gun → accelerating waveguide → bending magnet → target (or pass-through for electron mode) → flattening filter (or FFF configuration) → monitor chamber → collimation/MLC.
Key Variables	Waveguide type (standing vs. traveling wave); bending magnet angle (commonly 90°, 112.5°, or 270°); presence or absence of a flattening filter.
Clinical Application	These components collectively determine beam energy, flatness, symmetry, and achievable dose rate, all of which are verified during machine commissioning and routine QA.
Worked Example	Removing the flattening filter (FFF mode) substantially increases central-axis dose rate and produces a forward-peaked rather than flat beam profile, a design exploited in SBRT to shorten treatment delivery time.
High-Yield Pitfall	FFF beams have a fundamentally different (non-flat) dose profile that must be explicitly modeled in treatment planning — they cannot be planned identically to a flattened beam of the same nominal energy label.

Photon Interactions with Matter

Formula / Concept	Photoelectric effect dominates at low energy and scales with Z^3/E^3 ; Compton scatter dominates at intermediate (therapeutic) energies and depends on electron density rather than Z ; pair production requires photon energy above 1.02 MeV and increases with Z at high energy.
Key Variables	Photon energy E ; atomic number Z ; electron density of the medium.
Clinical Application	Explains why kV imaging beams show strong bone/soft-tissue contrast (photoelectric, strongly Z -dependent) while megavoltage therapy beams deposit dose relatively independent of atomic number (Compton-dominated), simplifying dose calculation across tissue types.
Worked Example	A therapeutic 6 MV beam passing through soft tissue and bone deposits dose according to electron density via Compton scatter, producing much less tissue-to-tissue contrast than a kV imaging beam would show for the same tissues.
High-Yield Pitfall	Assuming megavoltage dose deposition is strongly Z -dependent is a common error — at therapeutic energies, Compton scatter (electron-density dependent, not Z -dependent) is the dominant interaction, unlike the photoelectric-dominated kV imaging range.

Attenuation and Half-Value Layer

Formula / Concept	$I = I_0 e^{-\mu x}$; $HVL = 0.693/\mu$
Key Variables	μ = linear attenuation coefficient; x = absorber thickness.
Clinical Application	HVL is used clinically to characterize beam quality (effective energy) for kV beams and to design shielding barriers.
Worked Example	If a shielding material has an HVL of 1 cm for a given beam, then 3 cm of that material reduces beam intensity to $(1/2)^3$, or 12.5%, of its original value.
High-Yield Pitfall	The simple exponential attenuation model assumes narrow-beam ('good') geometry; broad-beam geometry includes a scatter contribution that causes measured attenuation to fall off less steeply than the pure exponential model predicts.

Electron Beams

Electron Beam Depth-Dose Characteristics

Formula / Concept	Electron percent depth dose rises to a maximum near the surface, falls off relatively sharply near the practical range (R_p), then trails off into a low-level bremsstrahlung tail.
Key Variables	R_p (practical range); R_{50} (depth of 50% dose, related to mean incident electron energy).
Clinical Application	Electron beams are chosen for treating superficial targets, such as skin lesions or chest wall boosts, where sparing of deeper underlying structures is desired.
Worked Example	A 9 MeV electron beam, with a practical range of roughly 4 to 5 cm, is a reasonable choice for treating a chest wall recurrence a few centimeters deep while sparing underlying lung and heart from significant dose.
High-Yield Pitfall	Selecting an electron energy that is higher than necessary for the target depth unnecessarily irradiates tissue deep to the target, since higher-energy electron beams have a correspondingly deeper range and dose falloff.

Electron Interactions: Collisional vs. Radiative Losses

Formula / Concept	Total mass stopping power is the sum of collisional (ionization/excitation) losses and radiative (bremsstrahlung) losses; radiative losses increase with both electron energy and absorber Z.
Key Variables	Electron kinetic energy; absorber atomic number Z.
Clinical Application	Explains the low-level bremsstrahlung 'tail' seen beyond the practical range on an electron PDD curve, and why higher-energy electron beams produce more x-ray contamination dose beyond the target.
Worked Example	A 20 MeV electron beam produces a more pronounced bremsstrahlung tail beyond its practical range than a 6 MeV beam, illustrating that radiative losses become proportionally more significant at higher electron energies.
High-Yield Pitfall	The x-ray contamination tail beyond an electron beam's practical range is a real, non-zero dose contribution to tissue beyond the intended target and should be accounted for, particularly at higher electron energies.

Dosimetry and Calibration

TG-51 Calibration Formalism

Formula / Concept	$D_w = M \times k_Q \times N_{D,w}^{Co-60}$ (conceptual form).
Key Variables	M = fully corrected chamber reading; k_Q = beam-quality conversion factor; $N_{D,w}$ = chamber calibration coefficient traceable to an accredited standards laboratory.
Clinical Application	TG-51 is the modern absolute dose calibration protocol used to ensure that a linear accelerator's monitor unit output corresponds accurately to a specific absorbed dose in water under reference conditions.
Worked Example	TG-51 establishes the absorbed dose per monitor unit under defined reference conditions; many clinics subsequently normalize machine output to approximately 1 cGy/MU at a specified reference geometry, but that normalization is a clinic choice, not a requirement dictated by TG-51 itself.
High-Yield Pitfall	TG-51 calibration is referenced to absorbed dose to water, a change from the older air-kerma-based TG-21 protocol — know that the modern calibration formalism is water-based, not air-based. Also do not confuse the protocol with a universal requirement that 1 MU equal exactly 1 cGy.

Ionization Chamber Physics (Bragg-Gray Cavity Theory)

Formula / Concept	Starting from dose to the cavity gas: $D_{medium} = D_{gas} \times (S/p)_{medium,gas}$. Starting from measured ionization: $D_{medium} = (Q/m) \times (W/e)_{gas} \times (S/p)_{medium,gas}$ (conceptual form).
Key Variables	Q/m = collected charge per unit gas mass; W/e = mean energy expended per unit charge of ionization in the gas (approximately 33.97 J/C for air); $(S/p)_{medium,gas}$ = mean mass collision stopping-power ratio between medium and gas.
Clinical Application	Provides the theoretical basis for relating ionization measured in a small gas-filled chamber cavity to absorbed dose in the surrounding medium, such as water or tissue.
Worked Example	An ionization chamber's air cavity does not significantly perturb the charged-particle fluence in the surrounding medium as long as the cavity is small relative to the range of the electrons crossing it, which is the core assumption underlying Bragg-Gray cavity theory.
High-Yield Pitfall	Do not multiply by W/e twice — if D_{gas} already represents absorbed dose in the gas, the W/e conversion has already been accounted for, and applying it again produces an incorrect result. The Bragg-Gray assumption can also break down at low energies or with larger cavities, requiring additional correction factors.

Percent Depth Dose, TAR, and TPR

Formula / Concept	$PDD = (D_d / D_{d_{\text{max}}}) \times 100$; $TAR = D_d / D_{\text{free-space}}$; $TPR = D_d / D_{\text{ref depth}}$.
Key Variables	Depth d ; SSD; field size; beam energy.
Clinical Application	PDD, TAR, and TPR are all used to characterize how dose changes with depth, but differ in their dependence on treatment distance, which determines which quantity is appropriate for isocentric versus SSD-based calculations.
Worked Example	A change in SSD changes PDD. This can be approximately estimated using the Mayneord F factor, although measured data or a more complete calculation is preferred when accuracy matters, since simply reading the same PDD table value at a different SSD without any correction introduces systematic error.
High-Yield Pitfall	PDD depends on SSD, while TAR and TPR do not — failing to recognize this distinction is a common source of error when converting between these quantities for different treatment setups. Also remember that the Mayneord F factor is only an approximation, principally accounting for inverse-square differences, and does not fully model changes in phantom scatter and field geometry.

Inverse Square Law

Formula / Concept	$I_1 / I_2 = (d_2 / d_1)^2$
Key Variables	d_1 , d_2 = distances from the source.
Clinical Application	Governs how dose rate changes with distance from a point source, relevant to both external beam SSD changes and brachytherapy source positioning.
Worked Example	For a free-space point-source example, if dose rate is 100 cGy/min at 80 cm source-to-point distance, at 100 cm it is approximately $100 \times (80/100)^2 = 64$ cGy/min.
High-Yield Pitfall	This idealized calculation applies to free-space intensity measured at defined source-to-point distances. In a patient or phantom, use source-to-calculation-point distance rather than SSD alone, and remember that scatter and finite-source effects can modify a purely inverse-square estimate.

Monitor Unit Calculation

Formula / Concept	$MU = \text{Prescribed dose} / (\text{cGy per MU at calibration} \times Sc \times Sp \times TPR \text{ or } PDD \times \text{wedge factor} \times \text{tray factor} \times \text{off-axis factor, as applicable})$.
Key Variables	Each factor is field-size, depth, and energy dependent and is measured during commissioning.
Clinical Application	This calculation converts a prescribed physical dose into the number of monitor units the accelerator must deliver, and is performed both by the treatment planning system and, traditionally, as an independent manual second check.
Worked Example	For a given prescription, field size, and depth, the reference calibration factor is adjusted by all relevant scatter, depth-dose, and modifier factors (wedge, tray, blocking) to solve for the monitor units required to deliver the intended dose to the specified point.
High-Yield Pitfall	Omitting any single relevant correction factor, such as a wedge or tray factor, is one of the most common sources of manual calculation error, which is why independent second-check calculations remain a standard radiation safety practice.

Treatment Planning

Beam Modifiers: Wedges and Compensators

Formula / Concept	Physical wedges are solid attenuators that produce a dose gradient across the field; dynamic (virtual) wedges achieve the same gradient by moving a collimator jaw during beam-on time; compensators correct for missing tissue or surface irregularity.
Key Variables	Wedge angle; wedge factor (ratio of wedged to open-field dose at a reference depth).
Clinical Application	Used to shape isodose distributions around curved surfaces or to compensate for tissue deficits, such as in breast or head and neck treatment.
Worked Example	A dynamic wedge achieves the same dose gradient as a physical wedge by moving one collimator jaw during treatment delivery, avoiding the beam-hardening effect that occurs when a physical wedge preferentially attenuates lower-energy photons.
High-Yield Pitfall	Physical wedges harden the beam (increase its average energy) by preferentially attenuating lower-energy photons, while dynamic wedges do not change beam quality — a distinction occasionally tested regarding beam-hardening effects.

Tissue Heterogeneity Corrections

Formula / Concept	Correction algorithms range in sophistication from simple effective path length methods to convolution/superposition and Monte Carlo methods that more accurately model electron transport near density interfaces.
Key Variables	Tissue density (from CT-derived electron density); algorithm type.
Clinical Application	Necessary for accurate dose calculation through or near lung, bone, and air cavities, where simple water-equivalent assumptions would misestimate dose.
Worked Example	A tumor located behind a large air cavity, such as in the paranasal sinuses, may receive a different dose than a homogeneous water-only calculation would predict, due to loss of electronic equilibrium near the air-tissue interface, an effect that becomes more pronounced at higher photon energies.
High-Yield Pitfall	Simpler correction algorithms, such as effective path length methods, do not accurately model dose near low-density interfaces such as lung-tumor boundaries the way convolution/superposition or Monte Carlo algorithms do — particularly relevant for small lung SBRT targets.

IMRT and VMAT Delivery

Formula / Concept	IMRT modulates beam fluence using multileaf collimator segments (step-and-shoot or sliding window) across multiple static gantry angles; VMAT delivers continuously modulated fluence during one or more gantry rotations by simultaneously varying MLC position, dose rate, and gantry speed.
Key Variables	Number of segments/control points; gantry arc range; dose rate modulation.
Clinical Application	Enables highly conformal dose distributions and improved organ-at-risk sparing compared with conventional 3D conformal techniques, at the cost of increased planning and QA complexity.
Worked Example	A VMAT plan for a head and neck target can achieve a highly conformal dose distribution while sparing the parotid glands, by modulating MLC position, dose rate, and gantry speed simultaneously during one or a few arc rotations, substantially reducing treatment time relative to static-field IMRT.
High-Yield Pitfall	The increased conformality of IMRT/VMAT increases sensitivity to daily setup uncertainty and interfractional anatomic change, making image guidance more clinically important than with less conformal 3D conformal techniques.

Dose-Volume Histograms and Plan Evaluation

Formula / Concept	A cumulative DVH plots the volume of a structure receiving at least a given dose, summarizing the 3D dose distribution within that structure as a 2D curve.
Key Variables	V_x (percent volume receiving at least dose x); D_x (minimum dose received by $x\%$ of the volume).
Clinical Application	Used to compare candidate treatment plans against established organ-at-risk dose constraints, such as lung V20 or spinal cord maximum dose.
Worked Example	A DVH showing that only 15% of total lung volume receives at least 20 Gy (a lung V20 of 15%) is generally considered an acceptably low risk of radiation pneumonitis under many clinical protocols, illustrating how DVH metrics translate into plan acceptability judgments.
High-Yield Pitfall	A DVH conveys no spatial information — two dose distributions with very different spatial patterns (such as a small hot region versus a larger area of moderate dose) can produce similar-looking DVH curves, so DVH review should always be paired with visual inspection of the dose distribution on the images.

Brachytherapy

TG-43 Dose Calculation Formalism

Formula / Concept	Dose rate is calculated from air kerma strength, the dose rate constant, a geometry function, a radial dose function, and an anisotropy function specific to the source model (conceptual form).
Key Variables	S_K (air kerma strength); Λ (dose rate constant); geometry, radial dose, and anisotropy functions.
Clinical Application	TG-43 is the standard modern formalism for brachytherapy treatment planning dose calculations, replacing older point-source or line-source approximations with measured or Monte Carlo-derived correction factors specific to each source design.
Worked Example	Two source models with identical air kerma strength can deliver different doses at the same distance if their dose rate constants differ, illustrating why source-specific TG-43 parameters, not air kerma strength alone, are required for accurate brachytherapy planning.
High-Yield Pitfall	Different physical source designs of the same isotope (differing in encapsulation or active length) require their own measured or simulated TG-43 parameter sets — do not assume dosimetric equivalence between sources of the same isotope from different manufacturers.

Brachytherapy Dose Falloff

Formula / Concept	Dose falls off approximately as $1/d^2$ near a brachytherapy source, dominating over other dosimetric effects at short distances.
Key Variables	Distance from the source; source geometry (point versus line source).
Clinical Application	Explains why organs immediately adjacent to an intracavitary applicator, such as the rectum or bladder in cervical brachytherapy, are disproportionately sensitive to small positioning changes.
Worked Example	Doubling the distance from a brachytherapy source reduces dose rate to roughly one-quarter of its original value, which is why millimeter-scale shifts in applicator or organ position near the source can meaningfully change the dose received by adjacent organs at risk.
High-Yield Pitfall	Unlike external beam therapy, where distance changes of a few centimeters have modest dosimetric impact, brachytherapy dose is extremely sensitive to millimeter-scale positioning changes near the source because of the steep inverse-square falloff at short range.

Imaging Physics Relevant to Radiation Oncology

CT Number and Electron Density Conversion

Formula / Concept	$HU = 1000 \times (\mu_{\text{tissue}} - \mu_{\text{water}}) / \mu_{\text{water}}$. The TPS uses a commissioned CT calibration to convert HU into the density and/or material information required by its dose-calculation algorithm.
Key Variables	Hounsfield Unit (HU); electron density and/or mass density; material assignment; CT energy/protocol.
Clinical Application	Essential for heterogeneity-corrected dose calculation in the treatment planning system. Traditional teaching emphasizes electron density as the key input, but modern algorithms may convert HU into electron density, mass density, and/or material composition depending on the algorithm — Monte Carlo systems in particular require material and density assignments.
Worked Example	A phantom with tissue-equivalent inserts of known density is scanned to establish the institution-specific CT-to-density/material relationship used by the TPS to convert each CT voxel's HU value into the information needed for accurate dose calculation.
High-Yield Pitfall	Do not assume every modern algorithm uses only electron density as its input — some require mass density and material composition as well. Metal artifacts on CT, such as from dental fillings or hip prostheses, produce erroneous HU values that can significantly distort dose calculation near the artifact unless specifically corrected.

MRI Basics for Radiation Oncology

Formula / Concept	T1- and T2-weighted sequences provide differing soft tissue contrast based on tissue relaxation properties; MRI uses no ionizing radiation and provides no direct electron density information.
Key Variables	T1/T2 relaxation times; sequence type; presence of contrast enhancement.
Clinical Application	Used for target delineation, particularly in CNS, prostate, gynecologic, and sarcoma radiotherapy, typically via MRI-CT fusion since MRI alone cannot be used directly for dose calculation without an electron density assignment.
Worked Example	A brain tumor may be poorly visualized on CT but clearly delineated on a T1-weighted post-contrast MRI, which is why MRI-CT fusion, rather than CT alone, is standard practice for target delineation in CNS radiotherapy.
High-Yield Pitfall	MRI images can have geometric distortion from gradient nonlinearity and magnetic field inhomogeneity that, if uncorrected, introduces spatial errors when fused with CT for treatment planning — an important and growing concern for MR-only planning workflows.

PET Physics and Standardized Uptake Value (SUV)

Formula / Concept	SUV is calculated from the measured activity concentration in tissue normalized to the injected dose and patient body weight (or other normalization factors, such as lean body mass).
Key Variables	Injected activity; uptake time; patient body weight or normalization metric.
Clinical Application	Used for staging, biologic target volume delineation, and treatment response assessment, most commonly with FDG as the radiotracer.
Worked Example	A pulmonary nodule with a markedly elevated SUV on FDG-PET raises suspicion for malignancy and may inform gross tumor volume definition for lung radiotherapy planning, though PET-based volumes require careful thresholding methodology to avoid over- or under-inclusion of the true tumor extent.
High-Yield Pitfall	SUV values are influenced by many technical factors, including blood glucose level, uptake time, scanner calibration, and patient body habitus — comparing SUV values across different scanners or protocols without standardization can be misleading.

Radiation Protection

Shielding and Barrier Design

Formula / Concept	Required barrier attenuation is calculated from workload (W), use factor (U), occupancy factor (T), and distance, typically expressed as a number of tenth-value layers (TVLs) of the shielding material.
Key Variables	Workload (patient throughput and dose per patient); use factor (fraction of workload directed at a given barrier); occupancy factor (fraction of time an area beyond the barrier is occupied).
Clinical Application	Used in vault design for linear accelerators and brachytherapy suites to keep occupational and public radiation exposure below regulatory limits.
Worked Example	A primary barrier, which must attenuate the full useful beam, generally requires substantially more shielding thickness than a secondary barrier, which only needs to attenuate leakage and scattered radiation — which is why vault walls are not uniformly thick in all directions.
High-Yield Pitfall	Primary and secondary barriers have very different shielding requirements based on the beam directions and workload they must attenuate; assuming uniform wall thickness is adequate for full vault shielding design is a common conceptual error.

Equivalent Dose and Effective Dose

Formula / Concept	Equivalent dose $H = D \times w_R$; effective dose $E = \sum(w_T \times H_T)$.
Key Variables	w_R = radiation weighting factor (photons/electrons = 1, alpha particles = 20); w_T = tissue weighting factor.
Clinical Application	Used to estimate stochastic risk from low-dose, whole-body or partial-body diagnostic and occupational exposures, and to compare risk across different radiation types and exposure patterns.
Worked Example	A diagnostic imaging exam with an effective dose of 10 mSv is used to estimate a theoretical whole-body stochastic risk, which is a fundamentally different concept from the physical dose delivered to a single organ during radiation therapy.
High-Yield Pitfall	Effective dose is primarily a radiation-protection quantity for estimating population-level stochastic risk and is not appropriate for quantifying or comparing therapeutic radiation doses. Radiotherapy can still produce stochastic late effects, particularly second malignancy risk — the reason effective dose is unsuitable here is its intended radiation-protection use, not an absence of stochastic risk from radiotherapy.

Quality Assurance and Statistics

Routine Machine QA Tolerances

Formula / Concept	Daily, monthly, and annual QA checks assess output constancy, beam energy, mechanical and radiation isocenter accuracy, MLC positioning, and imaging system alignment, each with an established tolerance appropriate to its clinical consequence.
Key Variables	Output constancy (percent deviation from baseline); mechanical tolerances (mm); tolerances generally tighten for high-precision techniques.
Clinical Application	Ensures that daily treatment delivery remains consistent with the parameters used during commissioning and treatment planning, catching drift or malfunction before it affects patient treatment.
Worked Example	A daily output check measuring 4% above baseline would exceed the commonly used $\pm 3\%$ daily output-constancy tolerance and should trigger action according to the clinic's QA policy before routine treatment continues — note that exactly +3% is at the tolerance boundary, not beyond it; only a value exceeding 3% is out of tolerance.
High-Yield Pitfall	Know the difference between a tolerance limit and an action level. QA tolerances are also not one-size-fits-all — they tighten considerably for high-precision techniques such as SRS/SBRT relative to conventional fractionated treatment, reflecting the reduced margin for error at small field sizes and high doses per fraction.

This reference summarizes widely used field-standard physics concepts for board review and is not a substitute for a full physics text or your institution's clinical calculation protocols. © Hasan Murshed, MD, MS — drhasanmurshed.com