EMPIRICAL

The Duality of Risk and Capacity in Vulnerability Management

MICHAEL ROYTMAN
CTO
EMPIRICAL SECURITY

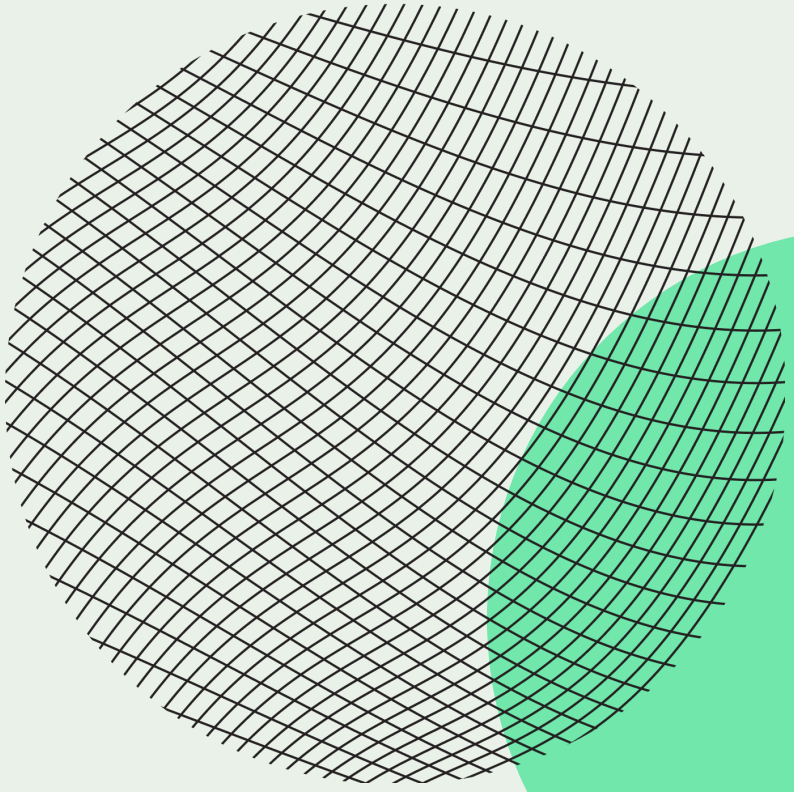

EMPIRICAL
OBSERVATIONS

Introduction

In most executive conversations about cybersecurity, the discussion centers on reducing the number of open vulnerabilities and staying below an agreed-upon risk threshold. These thresholds are often treated as fixed targets — a compliance benchmark or a comfort line that must not be crossed. The problem is that this view is static. It doesn't capture the dynamic relationship between the level of risk you're willing to tolerate and the operational capacity available to reduce it.

Optimization theory offers a more insightful lens. In that world, every problem has a dual: a mathematically linked counterpart that reframes the original question. In vulnerability management, the "primal" problem is how to minimize residual security risk with the capacity you have. The "dual" problem flips the perspective, asking: given a risk threshold, what level of remediation capacity is required to achieve it? The relationship between minimizing risk thresholds and their dual (maximizing remediation capacity) creates a framework for quantifying trade-offs, guiding investment decisions, and aligning operational work with strategic goals.

Thinking in terms of duality gives CISOs and CIOs a way to make resource allocation decisions based on measurable returns. If the marginal value of adding one more remediation sprint is a significant drop in residual risk, that's a clear case for adding budget or staff. If lowering the risk threshold by a small amount demands a steep increase in capacity, it may be better to accept a slightly higher level of residual risk and invest elsewhere. The model also works in the other direction: when vulnerability inflow spikes, such as after a major zero-day disclosure, you can calculate whether temporary capacity increases are worth the risk reduction they deliver.

Of course, any optimization framework is only as good as the data that feeds it. Many organizations unknowingly distort their models with stale or incomplete information. Vulnerabilities that were considered "hyper-critical" years ago may no longer pose meaningful re-exploitation risk, yet they still consume remediation resources. Insurance claims and breach reports can take more than a year to surface publicly, leaving models blind to current threat patterns. And just because a vulnerability hasn't been observed in the wild doesn't mean it's safe: it may simply be undetected. These blind spots shift both the primal and dual results, making it easy to invest capacity in the wrong places.

The challenge is compounded by the fact that security operates in what cognitive scientists call a wicked environment. Feedback is delayed, ambiguous, and often misleading. In such conditions, even experienced teams can lose calibration without deliberate measurement. For a duality model to work in practice, it needs robust feedback loops. On the primal side, this means validating that remediation actions are actually reducing

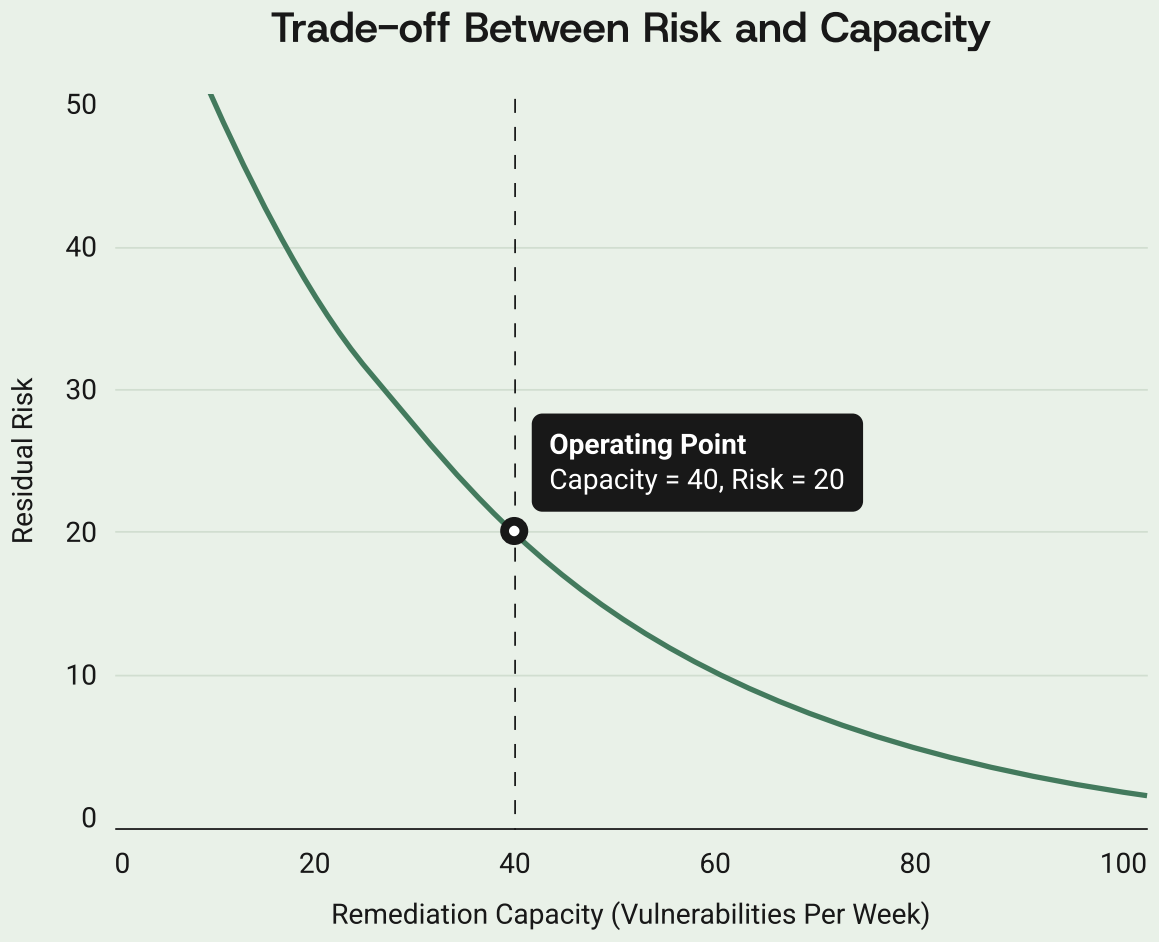

measurable risk, not just closing tickets. On the dual side, it means keeping the marginal value of capacity up to date as threats evolve. That requires correlating defensive actions with actual outcomes. For example, confirming that blocked attacks in the ids have no more vulnerabilities to target in the environment, rather than keeping detect and respond actions independent of preventative remediations and controls.

For CIOs, adopting this mindset can be transformative. First, it reframes vulnerability management from a reactive checklist to a balancing act between acceptable risk and operational bandwidth. Second, it provides a shared language for security and operations teams to discuss trade-offs with precision. And third, it allows executives to move beyond compliance-driven thresholds toward dynamic, data-informed decision-making.

By treating risk thresholds and remediation capacity as duals, leaders can measure exactly what each unit of additional capacity buys in risk reduction and what tightening or loosening thresholds will cost in operational terms. With accurate, timely data and strong feedback loops, this approach enables confident decisions about when to add capacity, when to adjust thresholds, and when to reallocate resources to prevention or detection. It's a way to ensure that every decision made in the boardroom is tied directly to measurable changes in security outcomes.

KEY TAKEAWAY

By treating risk thresholds and remediation capacity as duals, leaders can measure exactly what each unit of additional capacity buys in risk reduction and what tightening or loosening thresholds will cost in operational terms.

Making the Business Case for Remediation Capacity with Shadow Pricing

One of the most difficult questions facing cybersecurity leaders today is deceptively simple: What do we gain by increasing our remediation capacity?

One of the most difficult questions facing cybersecurity leaders today is deceptively simple: What do we gain by increasing our remediation capacity?

Every organization is constrained by some combination of remediation bandwidth, operational risk, and competing business demands. Yet the volume of detected vulnerabilities continues to grow, and the pressure to demonstrate progress never subsides. Against this backdrop, CISOs and CIOs are often forced to make difficult trade-offs between investing in more remediation capacity or accepting greater residual risk.

The traditional response to this problem has been to lean on best practices: prioritize critical vulnerabilities, follow vendor guidance, and apply industry benchmarks. But in increasingly complex and resource-constrained environments, best practices are no longer sufficient. They don't answer the question that security leaders need to ask: what do we gain by increasing our remediation capacity by one more unit? What is the actual value of fixing one more vulnerability?

This is where the concept of shadow pricing comes into play. In optimization theory, a shadow price quantifies the value of relaxing a constraint. In the context of vulnerability management, that constraint is remediation capacity. The shadow price tells us how much additional risk can be eliminated by patching one more vulnerability. This allows us to move beyond static rules and into a mode of dynamic, evidence-based prioritization.

Shadow pricing is especially powerful when applied to EPSS-based remediation strategies. Because EPSS scores are grounded in real-world exploitation likelihood, changes in threshold levels produce measurable changes in both remediation volume and risk reduction. This enables leaders to tie each unit of work to a specific marginal outcome—not just in terms of effort, but in terms of reduced exposure.

When the value of a remediation slot is known, it becomes possible to make smarter investment decisions. If an additional unit of capacity eliminates $800 of expected loss, but costs only $300 to create (whether through automation or staffing), the investment is self-justifying. If that same unit costs $1,200, the opportunity cost may be too high. Shadow pricing provides the missing economic rationale behind vulnerability prioritization decisions.

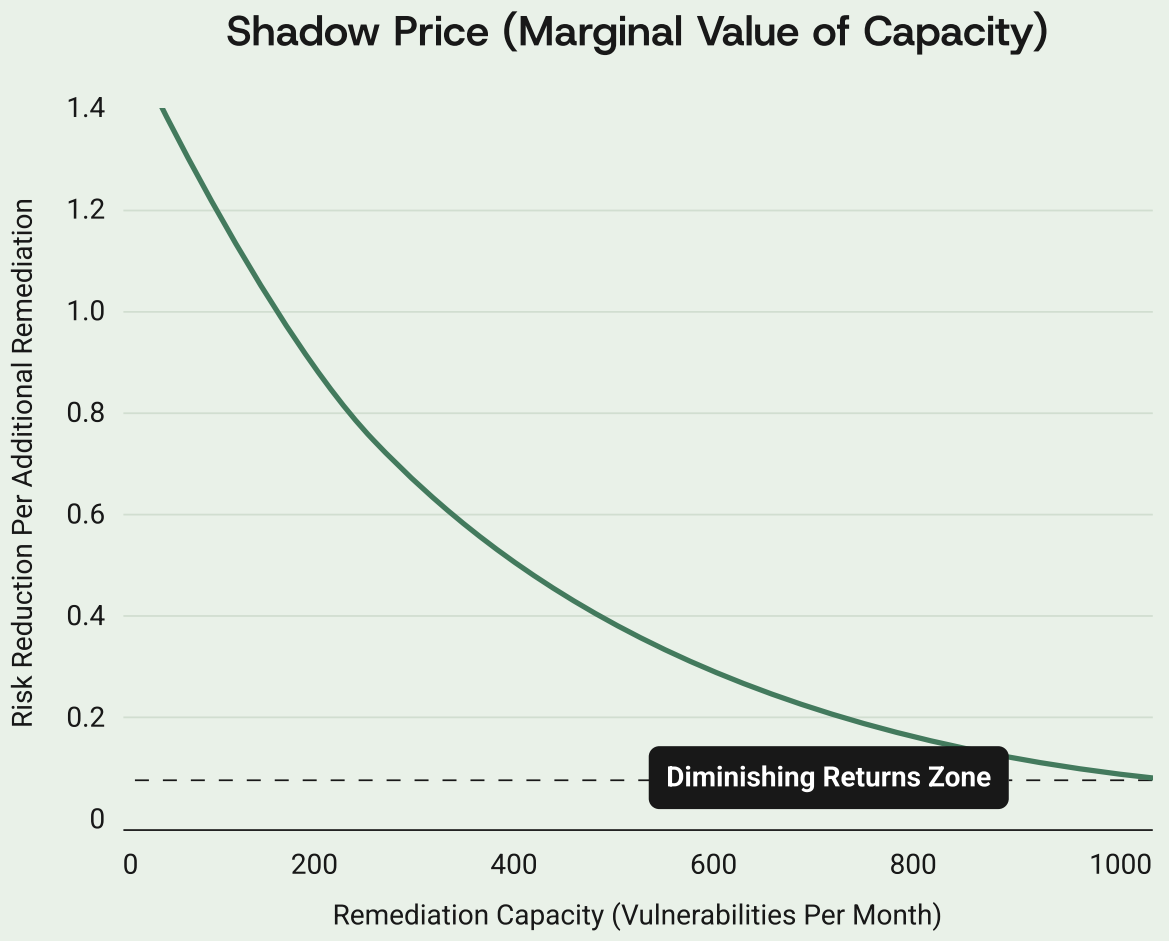

EMPIRICAL

Crucially, it also aligns security with finance. Security teams can present remediation plans and budget requests not as compliance obligations or staffing needs, but as investment decisions with measurable returns. This reframes the internal conversation: not “we need more people,” but “each additional person removes this much risk, worth this much in avoided loss.”

Perhaps most importantly, shadow pricing gives leaders a way to detect when capacity increases are no longer paying off. As thresholds lower and backlogs grow, the marginal value of additional remediation tends to decrease. At some point, adding more effort yields diminishing returns. Knowing where that point lies allows teams to shift resources away from remediation and into prevention, detection, or resilience without losing control of their risk posture.

The promise of shadow pricing is that it turns resource constraints into strategic levers. It provides a quantitative bridge between risk and capacity, and between security operations and business leadership. By adopting this mindset, CISOs and their teams can move from reactive patching to proactive optimization, making vulnerability management not just more efficient, but more defensible and more aligned with enterprise goals.

KEY TAKEAWAY

As thresholds lower and backlogs grow, the marginal value of additional remediation tends to decrease. At some point, adding more effort yields diminishing returns. Knowing where that point lies allows teams to shift resources away from remediation and into prevention, detection, or resilience without losing control of their risk posture.

How to Calculate a Shadow Price: Turning Model Thresholds into Operational Strategy

While shadow pricing is a powerful strategic concept, it becomes even more useful when applied directly to internal data. This post outlines a simple, concrete method for calculating the shadow price of remediation capacity using EPSS thresholds and remediation logs.

Let's assume your organization is currently operating with an EPSS threshold of 0.4, which means you prioritize remediation for all vulnerabilities with an EPSS score of 0.4 or higher. Based on this policy, your teams are remediating about 400 vulnerabilities per month, essentially operating at full capacity. Internal metrics estimate that this leaves your environment with a residual risk score of 20.

Now suppose you consider lowering the EPSS threshold to 0.3. This change would increase the number of vulnerabilities above the threshold to 500. However, the projected residual risk would fall to 12.0. That's an improvement of 8.0 units in risk reduction, requiring 100 more remediations.

To calculate the shadow price, divide the change in risk by the change in effort:

Risk reduction: 20 minus 12 equals 8 units

Additional remediations: 500 minus 400 equals 100

Shadow price: 8.0 divided by 100 equals 0.08 risk units reduced per remediation

If your organization assigns a financial value to each unit of residual risk then you can also express the shadow price in economic terms. Each additional remediation eliminates $800 in expected loss (0.08 multiplied by $10,000).

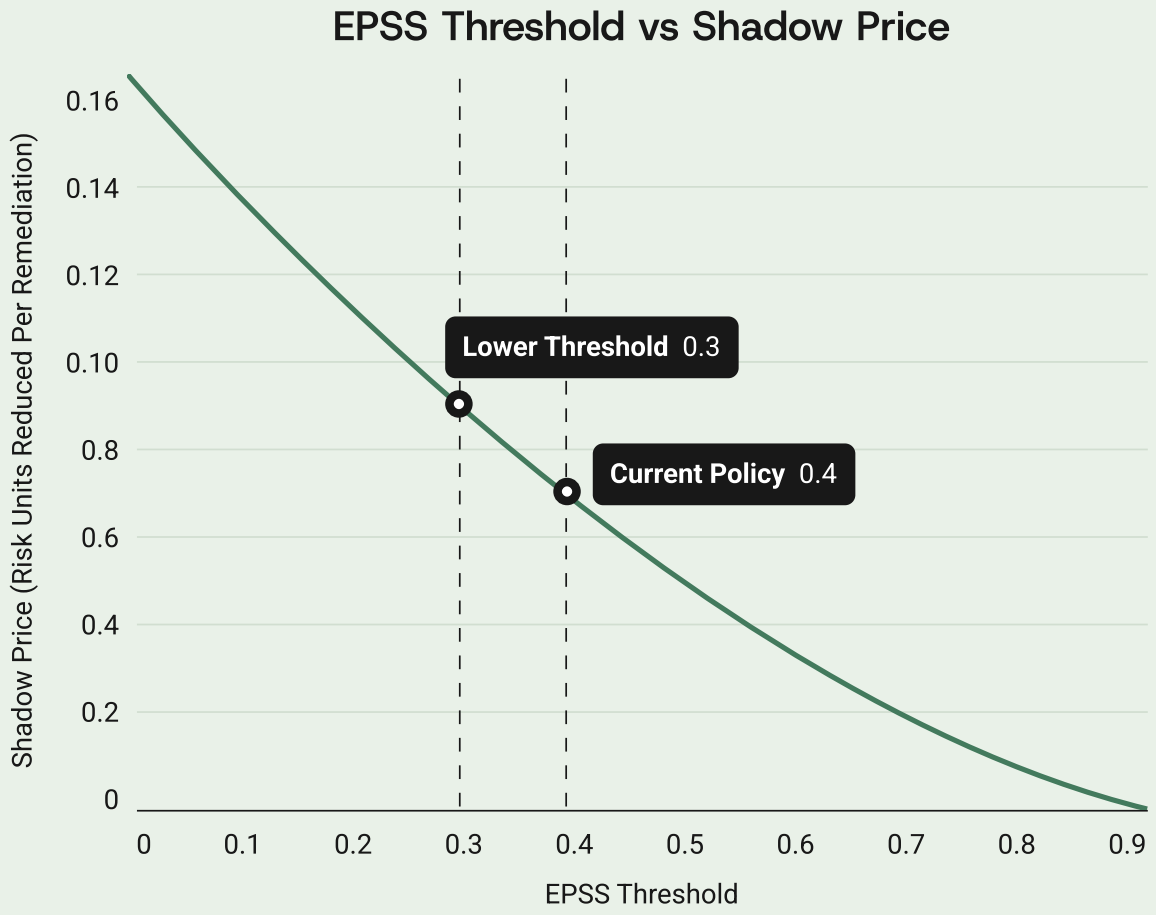

Now, compare that to the cost of creating one more remediation slot. If adding that capacity costs less than $800 per unit, the return is positive. If it costs more, it might be more efficient to hold the current threshold or reallocate resources.

This simple model transforms decision-making. Instead of setting EPSS thresholds arbitrarily or defensively, security and IT leaders can treat them as control parameters or adjustable levers in a real-time risk-reduction system. The shadow price tells you where the inflection point lies and when it makes sense to push harder versus hold steady.

Over time, this method becomes even more powerful as organizations collect better data. They can map how the shadow price changes under different threat conditions, staffing levels, or remediation tools. This enables adaptive thresholding based on real-world dynamics.

Shadow pricing does not require complex math or specialized software. It requires only curiosity, telemetry, and a willingness to measure trade-offs. By embedding this thinking into everyday security operations, teams can make smarter, faster, and more accountable decisions - not just about what to fix, but why it matters, and what each action is worth.

KEY TAKEAWAY

Shadow pricing does not require complex math or specialized software. It requires only curiosity, telemetry, and a willingness to measure trade-offs. By embedding this thinking into everyday security operations, teams can make smarter, faster, and more accountable decisions – not just about what to fix, but why it matters, and what each action is worth.

EMPIRICAL

Empirical builds data-driven models for security teams. We're building the world's first local models for cybersecurity, we maintain the world's most advanced global models, and we power existing open-source technology — EPSS, used by over 120 vendors today.

See how your model would differ — book a walkthrough
www.empiricalsecurity.com | info@empiricalsecurity.com

EMPIRICAL OBSERVATIONS