Heisenberg-Saturation Across Quantum Sensing Platforms: A Universal Law of Noise-Limited Precision

Deepinder Sidhu*
Department of Computer Science and Electrical Engineering
University of Maryland, Baltimore County, MD 21250, USA
sidhu@umbc.edu

July 18, 2026

Abstract

Quantum sensing platforms—including optical interferometers, atomic and ion clocks, NV-center magnetometers, and superconducting-qubit sensors—aspire to Heisenberg-limited precision, yet all exhibit an unavoidable breakdown of this scaling under realistic noise. Building on the Heisenberg-saturation framework introduced in earlier work, we present a *universal cross-platform law* that captures this breakdown in a single, architecture-independent form.

We decompose the total sensing resource into parallel and coherent components, distinguishing between the number of probes operated simultaneously and the depth of coherent interrogation. Across platforms, achievable precision is governed by the competition between ideal coherent amplification and noise-induced loss of state distinguishability. This competition is fully characterized by a platform-dependent coherence function, which encodes how rapidly noise erodes the information accumulated through deeper interrogation.

This formulation reveals that a single coherence parameter—such as an effective decay rate—is sufficient to predict three generic regimes common to all sensing technologies: an initial Heisenberg-like region, a coherence-limited optimum, and a saturation knee beyond which additional resources degrade rather than improve precision. Using literature-consistent coherence parameters, we present representative model curves for optical, atomic, superconducting, and NV-center platforms. Although coherence scales differ by orders of magnitude, all platforms obey the same qualitative law: **coherence—not particle number—sets the true precision frontier**. Heisenberg advantage is therefore not eliminated by noise, but strictly bounded by each platform's coherence budget.

The results presented here are model-based rather than experimental and are intended to provide a unified explanatory and comparative framework. Channel-specific precision-reduction laws and architecture-level resource partitioning are treated in companion papers; the present work isolates the cross-platform universality of noise-limited Heisenberg saturation.

1 Introduction

Heisenberg scaling,

$$\Delta\phi \propto \frac{1}{N}, \tag{1}$$

*This work is partially supported by Cyberspace Analytics Corporation.

has long served as the ideal benchmark for quantum-enhanced sensing and metrology [1]. In practice, however, every sensing platform—from squeezed-light interferometers and atomic clocks to NV centers and superconducting circuits—exhibits a systematic degradation away from this ideal behavior as resources increase. This degradation is not incidental; it reflects the finite coherence budget imposed by loss, dephasing, control error, and environmental coupling.

In Ref. [2], we introduced the concept of *Heisenberg-saturation* for a single platform: the emergence of a coherence-limited knee beyond which increasing resources no longer improves precision and may even degrade it. In a companion, channel-focused work [3], the same phenomenon was expressed as a *precision-reduction law* for sequential interrogation through a noisy quantum channel,

$$F_Q^{\mathrm{eff}}(L) = L^2\Lambda(L), \tag{2}$$

where L denotes the coherent interrogation depth and $\Lambda(L)$ captures channel-induced coherence loss.

The purpose of the present paper is to extend these ideas *across sensing technologies* and to address a cross-platform question that has remained largely implicit in the quantum-metrology literature:

How does Heisenberg-saturation manifest across different quantum sensing platforms, and what minimal set of parameters characterizes each platform's precision frontier?

We show that a single universal resource law,

$$F_Q^{\mathrm{eff}}(M, L) = (ML)^2\, C(L), \tag{3}$$

captures the precision behavior of all major quantum sensing platforms once the platform-specific coherence function $C(L)$ is specified. Here M denotes the number of probes used in parallel and L the coherent interrogation depth, so that the total resource is $N = ML$.

Here $C(L)$ is a dimensionless coherence-retention function that captures the cumulative degradation of phase coherence with increasing coherent interrogation depth L, with its form determined by the dominant noise mechanisms of the sensing platform.

Representative choices of $C(L)$, consistent with coherence windows reported in experimental and platform-level reviews of quantum sensing and metrology [**?**, **?**, **?**, **?**, **?**], yield model curves that place optical, atomic, superconducting, and solid-state spin sensors on a common resource axis.

Importantly, these curves are *model-based rather than experimental fits*. Their purpose is not to reproduce the performance of a specific device, but to expose a universal geometric structure underlying noise-limited precision across platforms. Despite large quantitative differences in coherence scales, all platforms obey the same qualitative law: Heisenberg-like scaling at small L, a coherence-determined optimum, and an inevitable saturation knee beyond which additional resources cease to improve precision.

1.1 Interpreting Precision: $\Delta\phi$ and the Normalized Metric $R(N)$

All quantum sensors considered here estimate a physical quantity by mapping it to an accumulated phase ϕ of a probe state. The achievable phase precision is therefore bounded by the effective quantum Fisher information,

$$\Delta\phi_{\mathrm{ach}}(M, L) = \frac{1}{\sqrt{F_Q^{\mathrm{eff}}(M, L)}}. \tag{4}$$

Although different platforms operate at distinct frequencies, wavelengths, or energy scales, they all reduce to estimating a phase shift, making $\Delta\phi$ a natural and universal metric for cross-platform comparison.

From quantum estimation theory, the quantum Cramér–Rao inequality gives

$$\mathrm{Var}(\hat{\phi}) \geq \frac{1}{F_Q^{\mathrm{eff}}(M,L)} \quad \Rightarrow \quad \Delta\phi_{\mathrm{ach}}(M,L) \equiv \sqrt{\mathrm{Var}(\hat{\phi})} \geq \frac{1}{\sqrt{F_Q^{\mathrm{eff}}(M,L)}}. \tag{5}$$

Throughout this work, $\Delta\phi_{\mathrm{ach}}$ should be interpreted as a *bound-to-achieved* precision: the bound is assumed achievable up to an $\mathcal{O}(1)$ efficiency factor common across platforms, which is suppressed because our focus is on scaling behavior and saturation knees rather than constant prefactors.

Interrogation time and coherent depth. The variable L counts coherent phase-imprinting steps (passes, cycles, or controlled interactions), not classical statistical repetitions. If each step has duration τ, the total sensing duration is $T = L\tau$. Equivalently, one may write $C(L) = C(T)$ and identify $\gamma L = (\gamma\tau)T$. Throughout this paper, either τ is fixed across comparisons or all time-dependence is absorbed consistently into the effective coherence function $C(L)$.

To compare platforms on a common, dimensionless scale, we introduce the normalized precision metric

$$R(N) \equiv \frac{\Delta\phi_{\mathrm{SQL}}(N)}{\Delta\phi_{\mathrm{ach}}(N)}, \tag{6}$$

where, by convention, $\Delta\phi_{\mathrm{SQL}}(N) = 1/\sqrt{N}$. Under this normalization,

$$R(N) = 1 \quad \text{(SQL)}, \qquad R(N) = \sqrt{N} \quad \text{(Heisenberg)}. \tag{7}$$

Because coherence loss depends on L rather than directly on N, the natural object is $R(M,L)$; a one-parameter curve $R(N)$ is obtained only after specifying how the total resource N is partitioned into $N = ML$ (e.g., by holding M fixed and increasing L).

1.2 Observation on Scale Factors and Repetitions

All expressions presented here refer to a single interrogation or architecture-level run. Numerical prefactors associated with generator normalization, estimator efficiency, and sensing duration are treated as $\mathcal{O}(1)$ when discussing scaling behavior.

An exception is the number of statistically independent repetitions ν_w. When ν_w repetitions are performed, the precision bound becomes

$$\Delta\phi \geq \frac{1}{\sqrt{\nu_w F_Q}}. \tag{8}$$

This factor rescales all precision curves uniformly and does not affect the location of saturation knees or platform-level optima. Accordingly, repetitions are omitted from plots and analytic comparisons and may be reinstated as a global prefactor when required.

2 Universal Resource Law: $N = ML$ and F_Q^{eff}

Following Ref. [2], we decompose the total sensing resource as

$$N = ML, \tag{9}$$

where M denotes the number of probes operated in parallel (photons, atoms, spins, qubits, etc.) and L denotes the coherent interrogation depth, such as multi-pass cycles, repeated phase imprinting, or an equivalent coherent interaction time.

In the absence of noise, coherent accumulation leads to ideal Heisenberg scaling,

$$F_Q^{\mathrm{ideal}}(M, L) \propto (ML)^2, \tag{10}$$

reflecting the quadratic growth of generator variance with the total coherent resource $N = ML$.

Under realistic operating conditions, however, noise limits the usable coherent depth and reduces the retained quantum Fisher information. We capture this effect through a platform-dependent *coherence function* $C(L)$ and define the effective, noise-limited quantum Fisher information as

$$F_Q^{\mathrm{eff}}(M, L) = (ML)^2 \, C(L). \tag{11}$$

The corresponding achievable phase precision is bounded by

$$\Delta\phi_{\mathrm{ach}}(M, L) = \frac{1}{\sqrt{F_Q^{\mathrm{eff}}(M, L)}}. \tag{12}$$

Equation (11) constitutes the central *universal resource law* of this paper. All platform-specific noise mechanisms enter exclusively through the single function $C(L)$, while the dependence on total resource factorizes universally as $(ML)^2$. This separation cleanly distinguishes between architectural resource allocation and coherence-limited performance.

A broad and analytically tractable class of coherence budgets is given by the exponential form

$$C(L) = e^{-\gamma L}, \tag{13}$$

where the effective decay rate γ compresses loss, dephasing, amplitude damping, and control noise into a single platform-dependent parameter. Under this model,

$$F_Q^{\mathrm{eff}}(M, L) = (ML)^2 e^{-\gamma L}, \qquad \Delta\phi_{\mathrm{ach}}(M, L) = \frac{e^{\gamma L/2}}{ML}. \tag{14}$$

The *Heisenberg-saturation knee* occurs at the coherence-limited optimum L^*, where $F_Q^{\mathrm{eff}}(M, L)$ ceases to increase with increasing L and $\Delta\phi_{\mathrm{ach}}(M, L)$ reaches its minimum. Beyond this point, additional coherent depth converts phase accumulation into decoherence rather than information gain, leading first to saturation and ultimately to degradation of precision. Importantly, the location of this knee is determined entirely by the coherence function $C(L)$ and is therefore platform-specific while remaining independent of architectural implementation.

Table 1: **Representative coherence parameters across sensing platforms.** Each γ is an effective decoherence rate used to generate model curves under the exponential coherence model $F_Q^{\mathrm{eff}}(M,L) = (ML)^2 e^{-\gamma L}$. Values are literature-consistent examples intended for comparative modeling, not experimental fits.

Platform	Representative γ	Qualitative coherence behavior
Optical interferometers	$\gamma_{\mathrm{opt}} \approx 0.02$	Low loss, long coherence window
Atomic / ion clocks	$\gamma_{\mathrm{at}} \approx 0.07$	Intermediate interrogation depth
Superconducting-qubit sensors	$\gamma_{\mathrm{sc}} \approx 0.25$	T_1/T_2-limited coherence
NV-center spin sensors	$\gamma_{\mathrm{NV}} \approx 0.60$	Rapid dephasing, strong spin bath

3 Representative Coherence Parameters Across Platforms

To illustrate how the universal resource law in Eq. (11) manifests across realistic sensing technologies, we assign representative coherence parameters γ to several major quantum sensing platforms. These values are chosen to be consistent with typical loss, dephasing, and coherence times reported in authoritative experimental and platform-level reviews of quantum sensing and metrology, and are *not* intended as fits to any specific experiment [**?**, **?**, **?**, **?**, **?**].

Rather, the purpose of these parameters is to generate *literature-informed model curves* that place disparate platforms on a common, dimensionless coherence axis. This approach exposes how rapidly each technology departs from ideal Heisenberg scaling as coherent interrogation depth increases, while preserving cross-platform comparability.

- **Optical interferometers (squeezed or entangled light).** Low per-pass optical loss and long coherence lengths motivate a small effective decoherence rate, for example $\gamma_{\mathrm{opt}} \approx 0.02$, consistent with analyses of loss-limited interferometry and quantum-enhanced optical sensing [4, 5].
- **Atomic and ion clocks.** Collective dephasing, laser phase noise, and environmental field fluctuations suggest an intermediate coherence window, for example $\gamma_{\mathrm{at}} \approx 0.07$, consistent with coherence times and stability limits reported for optical and ion-based clocks [**?**, **?**].
- **Superconducting-qubit sensors.** Finite energy relaxation and dephasing times (T_1 and T_2) imply a comparatively short coherence depth, for example $\gamma_{\mathrm{sc}} \approx 0.25$, consistent with coherence budgets in contemporary superconducting-qubit architectures [**?**, **?**].
- **NV-center and solid-state spin sensors.** Strong coupling to surrounding spin baths, inhomogeneous broadening, and surface noise produce rapid dephasing, for example $\gamma_{\mathrm{NV}} \approx 0.60$, consistent with experimentally reported coherence limits in solid-state spin sensing [**?**, **?**].

4 Cross-Platform Heisenberg-Saturation Curves

Effective QFI versus interrogation depth

Figure 1 shows representative model curves of the effective quantum Fisher information

$$F_Q^{\mathrm{eff}}(M,L) = (ML)^2 C(L), \tag{15}$$

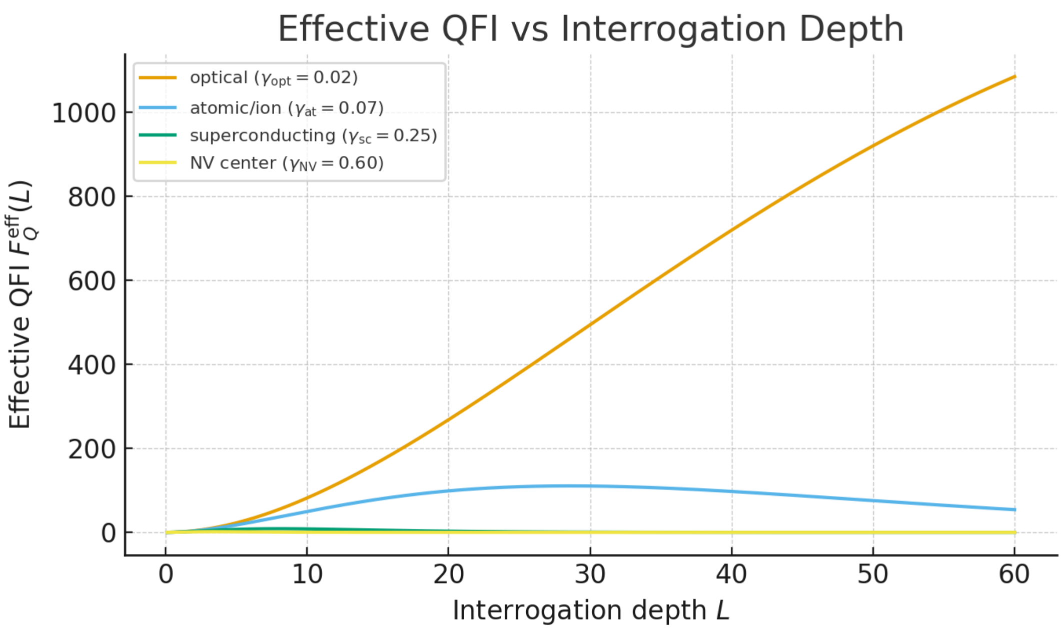

Figure 1: **Effective QFI across sensing platforms.** Model curves of $F_Q^{\text{eff}}(M, L) = (ML)^2 e^{-\gamma L}$ for representative platforms: optical interferometers ($\gamma_{\text{opt}} = 0.02$), atomic/ion clocks ($\gamma_{\text{at}} = 0.07$), superconducting-qubit sensors ($\gamma_{\text{sc}} = 0.25$), and NV-center solid-state sensors ($\gamma_{\text{NV}} = 0.60$). Smaller γ sustains coherent information buildup to larger L, whereas larger γ produces earlier saturation. Curves are literature-informed models rather than experimental data.

as a function of coherent interrogation depth L for four major sensing platforms, using the exponential coherence model $C(L) = e^{-\gamma L}$ introduced in Eq. (13).

For visual comparability, the number of parallel probes M is held fixed across all platforms. Under this normalization, the curves differ only through the platform-dependent coherence parameter γ. Changing M rescales all curves vertically by a factor of M^2 and therefore does not affect the location of the saturation knee.

In all cases, $F_Q^{\text{eff}}(M, L)$ initially increases quadratically with L, reflecting ideal Heisenberg scaling under coherent accumulation. As L increases further, coherence loss progressively suppresses the retained information, causing each curve to reach a maximum and then decrease. The location of this maximum defines a platform-specific saturation depth L^*, determined entirely by the coherence budget encoded in $C(L)$.

Achievable precision versus interrogation depth

The same physics is revealed from the complementary perspective of achievable phase precision. Using

$$\Delta\phi_{\text{ach}}(M, L) = \frac{1}{\sqrt{F_Q^{\text{eff}}(M, L)}} = \frac{e^{\gamma L/2}}{ML}, \tag{16}$$

Fig. 2 shows the corresponding precision curves on a semi-logarithmic scale.

Taken together, Figs. 1 and 2 directly answer the cross-platform question posed in the Introduc-

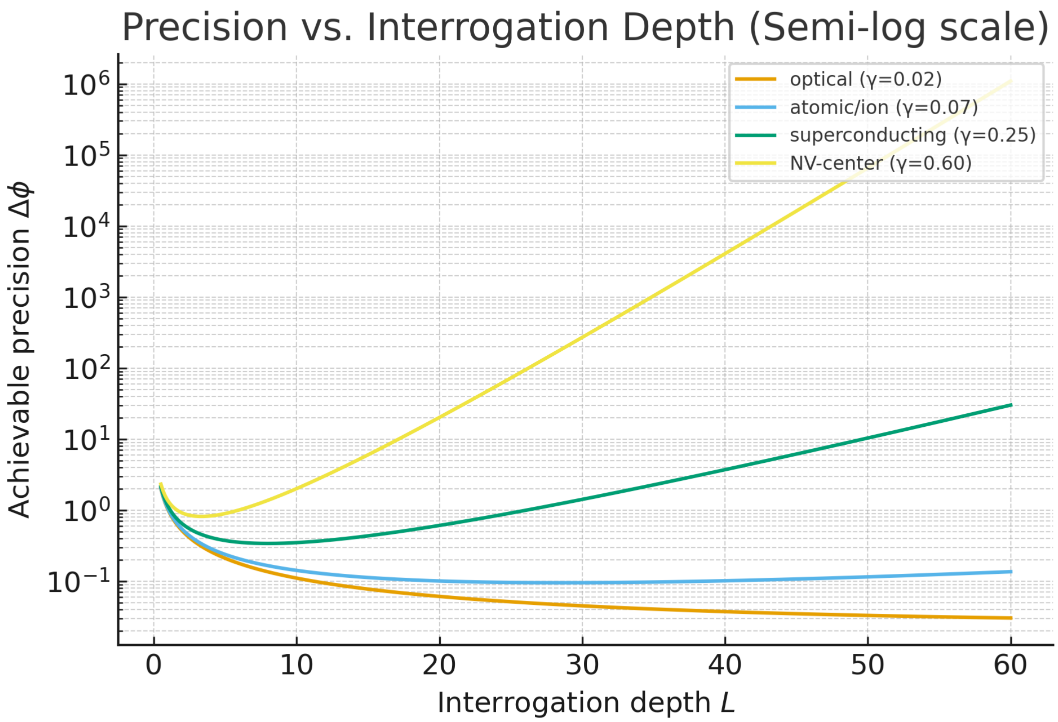

Figure 2: **Achievable precision versus interrogation depth across sensing platforms.** Phase uncertainty $\Delta\phi_{\rm ach}(M, L) = e^{\gamma L/2}/(ML)$ (semi-logarithmic vertical scale) for the same representative platforms shown in Fig. 1. Each platform exhibits an initial Heisenberg-like improvement with increasing L, followed by a coherence-limited minimum and eventual degradation. Smaller γ shifts the optimum to larger L and lowers the attainable precision floor, while larger γ produces earlier saturation and a more rapid loss of quantum advantage.

tion: *how far from Heisenberg scaling must each platform fall?* The answer is encoded entirely in the coherence function $C(L)$—or, under the exponential model, in the single effective parameter γ that characterizes each platform's coherence budget.

5 Implications for Cross-Platform Sensor Design

The universal resource law in Eq. (15), together with the cross-platform behavior illustrated in Figs. 1–2, leads to a set of general design principles that apply uniformly across quantum sensing technologies:

- **Coherence, not particle number, sets the precision frontier.** Increasing either the number of probes M or the interrogation depth L is beneficial only while the product $(ML)^2C(L)$ continues to grow. Beyond the saturation knee L^*, additional coherent depth yields no net information gain and ultimately degrades performance.
- **Each platform has a characteristic saturation depth.** The functional form of the coherence function $C(L)$—and, in particular, the effective decay parameter γ—determines both the location of the optimal interrogation depth L^* and the maximum achievable quantum Fisher information.

- **Coherence engineering dominates brute-force scaling.** Reducing γ through improved materials, environmental isolation, control fidelity, dynamical decoupling, or noise mitigation shifts the saturation knee to larger L and lowers the achievable precision floor far more efficiently than simply increasing the total resource N.
- **Resource allocation must respect the coherence budget.** Architectures that concentrate resources near the coherence-matched region around L^* preserve quantum advantage more robustly than strategies that push interrogation depth beyond the noise-limited window.

In short, the relevant design question is not *how large can N be?* but rather *how much of N can remain coherent under the platform's noise profile?*

5.1 Technology Trajectories (2025–2030): Where Can the Platforms Go?

Because the effective quantum Fisher information,

$$F_Q^{\mathrm{eff}}(M, L) = (ML)^2 C(L), \tag{17}$$

depends on platform-specific details only through the coherence function $C(L)$ (or an effective parameter γ), near-term technology progress can be visualized geometrically as a systematic deformation of the curves shown in Figs. 1 and 2.

- **Optical interferometers.** Further reductions in optical loss lower γ_{opt}, pushing the saturation knee to larger L and enabling deeper coherent buildup.
- **Atomic and ion clocks.** Improved magnetic-field control, collective noise suppression, and advanced interrogation protocols reduce γ_{at}, extending the usable coherent interrogation window.
- **Superconducting-qubit sensors.** Improvements in coherence times T_1 and T_2 directly reduce γ_{sc}, shifting the optimal operating point toward larger coherent depth.
- **NV-center and solid-state spin sensors.** Isotopic purification, surface engineering, and optimized dynamical decoupling can reduce γ_{NV} and may also modify the functional form of $C(L)$, extending the near-Heisenberg regime.

In every case, investment in coherence produces a simple geometric effect: **the curves move rightward and downward**, saturation knees shift to larger L, and each platform operates closer to the Heisenberg ideal over a broader region of parameter space.

6 Conclusion

Heisenberg-saturation was originally identified in Ref. [2] as a single-platform phenomenon: increasing coherent resources eventually consumes coherence faster than it improves precision. Here we have shown that the same mechanism applies uniformly across quantum sensing platforms once the total resource is expressed as

$$N = ML, \tag{18}$$

and the effective quantum Fisher information takes the universal form

$$F_Q^{\text{eff}}(M,L) = (ML)^2 C(L). \tag{19}$$

Using representative, literature-consistent coherence parameters for optical, atomic, superconducting, and NV-center platforms, we constructed model curves that illustrate how each technology approaches—and ultimately departs from—ideal Heisenberg scaling. The central conclusion is universal:

Each platform falls from Heisenberg exactly as fast as its coherence budget $C(L)$ dictates.

The true precision frontier is therefore not set by particle number alone, but by the ability to preserve coherence over the interrogation depth required for sensing. This cross-platform law provides a concrete and unifying design rule for future quantum sensors: **engineer coherence first, and only then scale resources**.

References

[1] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum metrology. *Physical Review Letters*, 96(1):010401, 2006.

[2] Deepinder Sidhu. Heisenberg-saturation across quantum sensing platforms: A universal law of noise-limited precision, 2025. Condensed conference manuscript.

[3] Deepinder Sidhu. Precision reduction in quantum channels: How far from heisenberg must we fall?, 2025. Condensed conference manuscript.

[4] V. Giovannetti, S. Lloyd, and L. Maccone. Advances in quantum metrology. *Nature Photonics*, 5:222–229, 2011.

[5] Rafał Demkowicz-Dobrzański, Marcin Jarzyna, and Jan Kołodyński. Quantum limits in optical interferometry. In Emil Wolf, editor, *Progress in Optics*, volume 60, pages 345–435. Elsevier, 2015.