

Quantum beneath the Ground: Commercial Tests for Subsurface Sensing

A Field-Evidence and Capital-Allocation Framework for Gravity-Based Subsurface Decisions

Authored by Chennakeshav Adya

Independent Researcher

October 2026

ABSTRACT

Quantum gravity instruments measure a physical signal that can help constrain an underground decision. A published field demonstration detected a buried tunnel with a cold-atom gravity gradiometer. That result establishes a specific measurement capability under tested conditions. It does not establish the location of ore, the productivity of an aquifer or geothermal reservoir, or the containment of injected carbon dioxide. This paper develops a commercial test for buying, deploying or financing quantum subsurface sensing. It asks whether a sensor changes an actionable decision after comparison with available geophysical methods, and whether that change is worth its complete cost.

The paper separates instrument performance, geophysical interpretation, field confirmation and transaction value. It proposes a physical-property map, a site-specific comparator, an independent blind-survey protocol, a rights and data ledger, and payment gates tied to decision evidence. Water, minerals, geothermal resources and geological storage present different targets and different truth tests. A fictional six-site programme illustrates the value of an avoided unsuccessful test after including survey cost and the risk of a missed viable site. Every probability and cash input in that example is an author assumption, not an observed sensor result or market benchmark.

The commercial proposition is strongest where a measured gravity difference resolves a costly choice that existing data leave open. The buyer should retain the ability to stop after a failed field test, obtain the data and interpretation needed to reproduce the result, and commission an independent geophysical and domain-specific review. The paper provides a route from a bounded pilot to a scale decision; it does not recommend a particular sensor vendor, asset or security.

JEL Classification: G31, G32, L72, O32, Q25, Q42

Keywords: quantum sensing, atom interferometry, gravity gradiometry, subsurface imaging, groundwater, geothermal, mineral exploration, carbon storage, value of information

1 Define the decision before the instrument

An underground asset can consume substantial capital before anyone sees the property on which its value depends. A mineral investor may have to select a drill target. A water utility may need to locate an aquifer boundary or measure storage change. A geothermal developer may need to decide where to place a confirmation well. A carbon-storage operator may need to distinguish an injection response from background variation. These are different decisions with different costs, deadlines and acceptable uncertainty. The first question for a proposed sensor purchase is which of them changes when the new measurement arrives.

A quantum instrument does not create a direct picture of every underground object. Cold-atom gravimeters use the interference of matter waves to measure gravity; a gradiometer measures a difference in gravity between two positions. The resulting signal reflects the distribution and change of mass around the instrument. The US Geological Survey explains that gravity methods infer subsurface density contrasts and are commonly interpreted with other geophysical measurements [3]. The purchaser therefore needs a geological or engineering hypothesis before commissioning a survey. If two plausible underground configurations generate similar signals, a more sensitive instrument may still leave the capital decision unresolved.

The investment committee should write a one-page decision statement. It should name the site, the right to enter and survey it, the competing physical models, the next expensive action and the evidence that would change that action. It should also state the conventional alternative available now. A technology supplier can then tender a test against the same decision rather than an unbounded promise to map the underground. The contract and valuation should follow the evidence required for that decision.

2 Establish what has actually been demonstrated

In 2022, researchers reported a field survey with a cold-atom gravity gradiometer that detected a buried tunnel beneath a road [1]. The experiment used a known tunnel and a site model incorporating the surrounding terrain and structures. The authors reported a gravity-gradient anomaly consistent with the target and discussed the uncertainty in the inferred position and depth. The experiment is important because it moved quantum gravity sensing into a real outdoor setting. Its target, geometry, survey conditions and independent site knowledge define what the result supports.

The same paper reports that the prototype's repeatability during the survey was similar to commercial gravimeters and that systematic effects remained [1]. A buyer should read that qualification beside any claim of laboratory sensitivity. A sensor can have a strong statistical measurement and still need better control of alignment, vibration, survey speed, positioning, calibration or interpretation. An outdoor demonstration of a shallow tunnel is not a measured success rate for minerals, groundwater, deep geothermal reservoirs or carbon storage. Those applications require their own physical forward models and field comparisons.

The commercial evidence file should distinguish four statements: the instrument produced a reproducible signal; the signal was interpreted as a particular subsurface feature; that interpretation was confirmed independently; and the confirmation improved an economic decision. A vendor may be able to prove the first without proving the fourth. The buyer should pay for each stage according to the deliverable actually received and the risk still carried by the buyer.

3 Match the sensor to the physical property

Gravity is responsive to mass. A density contrast can arise from a void, a change in rock type, a fluid redistribution, an orebody or human infrastructure. The direction and size of the signal depend on geometry, depth, density contrast and surroundings. A single gravity anomaly has no unique material label. The USGS groundwater-geophysics manual highlights the non-uniqueness of geophysical interpretation [4]. A mineral target cannot therefore be identified as a specific commodity from a gravity map alone. Sampling, drilling and other measurements remain central.

Quantum sensing describes how the instrument measures, rather than a new law connecting signal to resource. Atom interferometry can improve particular gravity measurements, but the value of an improvement depends on the survey environment and decision. A field trial should compare the uncertainty in the final interpreted target under two complete workflows, not merely the noise floor of two instruments. Each workflow includes survey layout, positioning, environmental corrections, processing, geophysical inversion and reviewer judgement. The conventional comparator may already be adequate when the expected target produces a large signal.

An investor should ask for a forward model before field acquisition. The model should specify the range of target depths, dimensions and density contrasts that could matter. It should also simulate plausible confounders: topography, nearby buildings, buried services, weather-related water changes and known geological layers. If the predicted difference between a viable target and a confounder is below total survey uncertainty, the project needs another measurement or a different physical method. Paying more for the same ambiguity is unlikely to resolve the investment choice.

4 Choose a use case with a real truth test

The most bankable early use case is one where the target and its counterfactual can be independently checked within a funded pilot. A shallow void or known utility corridor can be mapped against drawings, excavation or another survey. A water-storage application can compare repeated gravity observations with wells, pumping records and rainfall or recharge data. A geothermal application can compare time-varying mass observations with production and injection records. A mineral campaign can test a small set of interpreted structures through drilling and assays. Each test demands different permissions, seasonality and expertise.

The Birmingham tunnel demonstration provides an engineering benchmark for a particular shallow feature [1, 2]. USGS hydrologic work demonstrates that repeat gravity measurements can inform groundwater-storage change, although those cited programmes use established gravity methods rather than proving a quantum advantage [5]. GFZ's Equip-G programme plans quantum-gravimetry campaigns in a geothermal field to study mass redistribution [6]. DOE's NETL lists high-precision gravity among the methods relevant to carbon-storage

monitoring [7]. Together, these sources support a set of plausible test designs. They do not establish a single sensor's commercial performance across all four domains.

A sensible portfolio order begins with a site where ground truth is accessible, then tests transfer to a second geology and decision. The first site is an engineering and measurement trial. The second tests whether the claimed advantage survives a different context. A third, independently procured project can test whether customers will pay a price that covers deployment and support. Bundling these steps into one headline pilot risks confusing an instrument demonstration with product-market evidence.

Table 1. Different underground questions require different evidence

Decision	Gravity-relevant observable	Independent truth test	Essential comparator
Shallow void or service	Local density contrast	Drawings, excavation or established survey	Conventional gravimetry and utility mapping
Aquifer storage change	Change in water mass over time	Wells and water-balance observations	Repeat conventional gravity network
Mineral target	Structure or density contrast	Drilling, assays and geology	Existing gravity, magnetic and seismic data
Geothermal reservoir	Mass redistribution with operation	Wells, temperatures and flow records	Existing monitoring suite
Carbon storage	Change in subsurface mass distribution	Injection records and approved monitoring	Seismic, wells and high-precision gravity

Author-developed comparison based on the cited physical methods; no row establishes a quantum performance advantage.

5 Build a site-specific conventional comparator

The comparison is a procurement decision, so the baseline must represent the method the customer would otherwise buy. For a buried utility, that may be records, ground-penetrating radar, electromagnetic methods, conventional gravity and selective potholing. For groundwater, it may be wells and a repeat gravity network. For minerals, it may be geological mapping, magnetic or electromagnetic data, established gravity and a limited drilling programme. For carbon storage, it may be an approved monitoring plan built around seismic, wells, pressure and geochemistry. The baseline must have a budget and acceptance criterion.

NETL lists several methods used to monitor geological carbon storage, including well logging, fluid sampling, seismic imaging, high-precision gravity and electrical techniques [7]. It also warns that many methods detect an interpreted property change rather than carbon dioxide directly. This is a useful procurement discipline. The quantum sensor should be compared with the least expensive credible package that resolves the same uncertainty. A result may be valuable as a complement to that package even when it cannot replace it. The contract should specify the role the sensor claims to play.

The evaluator should blind the interpretation where practicable. Both teams receive the same site access, target definition, known background data and time window. They submit predictions before the independent truth test. Report the number of sites surveyed, time spent, incomplete stations, false positives, missed targets and unresolved cases. Include remediation, mobilisation, safety and data-processing costs. A comparison of instrument specification sheets alone would omit the expenses and errors that determine customer value.

6 Treat the underground image as an inverse problem


An inversion estimates an underground configuration from measurements at the surface or in a borehole. Different combinations of depth, geometry and density can fit the same gravity observation. The published tunnel study explicitly describes ambiguity between model parameters, even with a known target and auxiliary topographic information [1]. The USGS describes similar limits in gravity-based interpretation of basin geometry [3]. An image produced from an inversion is therefore a conditional model, not a photograph.

Every commercial interpretation should state the assumptions that constrain the inversion. These may include borehole control, geological maps, topography, known building foundations, instrument height, density measurements or a prior distribution. The report should show how the answer changes when those assumptions

move. A colourful slice through one preferred model can make a target appear more certain than the data warrant. The buyer should receive alternative models that fit the observations and a description of what independent test would distinguish them.

The relevant output for a capital decision is a bounded range of target locations or changes, linked to the cost of acting on an error. If two plausible inversions place a feature on opposite sides of a proposed drill pad, the drill cannot be approved merely because the plots look similar. A limited confirmatory borehole or another geophysical measurement may be cheaper than buying a denser gravity survey. The committee should compare the expected cost of that next information step with the cost of the wrong irreversible action.

Figure 1. Author-developed chain from physical measurement to capital decision



The chain separates observed signals from interpreted targets and independently checked outcomes.

7 Separate absolute gravity from gradients and time change

An absolute gravimeter reports local gravity against a reference. A gradiometer measures how gravity differs between positions. A time-lapse programme compares observations made at different times. These are related measurements with different error budgets. A shallow void may yield a local spatial anomaly. Monitoring fluid movement in an operating reservoir requires repeat measurements and correction for changes unrelated to the reservoir. The sensor architecture and survey plan should follow the target signal, not the prestige of a technology label.

The USGS describes repeat gravity monitoring as a way to estimate groundwater-storage change and notes the need to consider atmospheric, tidal and elevation effects [5]. GFZ's geothermal programme plans an absolute quantum gravimeter alongside superconducting gravimeters to provide reference and time-series information [6]. Those projects show why positioning, station stability and repeated calibration matter. A quantum sensor with a favourable single-station uncertainty may still deliver poor time-lapse interpretation if stations move, ground subsides or a local hydrologic change is ignored.

Procurement documents should therefore separate sensor noise, total survey repeatability and uncertainty in the interpreted quantity. The supplier should state which metric has been measured at the buyer's site and which comes from a laboratory or a different geology. A project that needs an answer within days may place high value on deployment speed. A project monitoring over years may care more about drift, maintainability and continuity when equipment is replaced. These are commercial performance attributes that cannot be reduced to one sensitivity number.

8 Design the first blind survey

A credible pilot starts with a finite area and a register of candidate features. The customer should identify safe survey access, utility records, known structures, benchmark sites and environmental constraints. The vendor should document its instrument configuration and freeze the processing method before receiving the answer key. An independent surveyor should record coordinates and elevations. An independent geophysicist should review corrections and inversion. The domain owner should define the cost of a wrong decision before results are unblinded.

Blindness has limits in a real site. A visible road alignment or wellhead can suggest where to look. The protocol should state which facts are available to both the conventional and quantum teams. If the vendor is allowed to use privileged maps, the same maps should be given to the comparator, or the claimed incremental value should include the price of those maps. Repeating the survey on an undisclosed line can test whether the anomaly reproduces. A favourable result at one known point should not substitute for performance over the predefined evaluation set.

The acceptance report should include station-level data, rejected observations, environmental logs, processing code or a reproducible specification, and all interpretations submitted before truth was revealed. It should report failures and ambiguous sites. The first pilot can confirm that the instrument works in the buyer's environment. The next purchase decision concerns whether its output saves money or avoids loss at the scale the buyer intends to deploy. Those are different gates with different evidence.

9 Price the option to learn before pricing a resource

The strongest commercial claim is often a reduction in uncertainty before a costly drill, excavation or injection decision. The buyer may pay for a sensor pilot as an option to learn, with the right to stop if it does not change the decision. That option can be worth more than the sensor's hardware value and less than the value of the resource that might later be confirmed. The price must reflect the incremental decision benefit after the cost of acquiring and interpreting the new measurement.

Consider a project where the committee must choose between drilling two sites and waiting for more evidence. The conventional geophysical programme may already rank them. If the quantum survey only repeats that ordering, its incremental value can be limited even when it is technically excellent. If it changes which site is drilled and an independent truth test later supports that change, the buyer can evaluate avoided work or protected option value. A change in ranking is not sufficient by itself. The supplier must show that the change was correct under the prespecified endpoint.

A contract can price delivery, validation and decision benefit separately. A modest initial payment can fund a bounded survey and reproducible data package. A later payment can depend on an independently reviewed accuracy or time-to-decision threshold. A resource-linked royalty requires special care because the resource, rights and development work may be supplied principally by other parties. Transaction counsel should define the instrument provider's claim so that later financing is not unexpectedly burdened by an open-ended contingent payment.

10 Keep the four applications economically distinct

Water, minerals, geothermal and storage are often grouped as subsurface markets. Their paying customers, regulatory duties and value units differ. A water manager may value a better estimate of storage change or a more efficient monitoring network. A mineral explorer values a target choice and the right to test it. A geothermal developer values avoided dry drilling, reservoir knowledge and later flow. A storage operator values monitoring that satisfies a site-specific plan and supports safe operation. The sensor is relevant only where its measured property bears on that unit of value.

The USGS has used gravity methods to estimate groundwater-storage change [5]. That is evidence for gravity as a hydrologic method, not evidence that every aquifer customer needs a quantum instrument. Mineral exploration already uses gravity with magnetic and other methods to map structures and density contrasts [3, 8]. GFZ's project explicitly investigates quantum measurements in an operating geothermal setting [6]. NETL treats high-precision gravity as one method among several for carbon-storage monitoring [7]. These distinct source contexts should inform separate bids, budgets and field protocols.

Selling one platform across all four requires evidence of transfer. Equipment may be portable, while interpretation software, station spacing, corrections and truth tests remain specialised. A buyer of the sensor company should value proven contracted revenue in each application separately. Prospective cross-selling belongs in a scenario until customers have run and paid for comparable deployments. Technical capability is a necessary input to a mandate; it is not a substitute for budget authority and procurement evidence.

11 Test the water case against wells and budget decisions

For groundwater, the commercially relevant question may be whether repeat gravity measurements improve a water-budget estimate enough to change abstraction, recharge or infrastructure decisions. The signal is a change in mass over time, not a direct measure of hydraulic head or water quality. The USGS notes that gravity observations can be combined with water levels to estimate aquifer properties and that a network can inform regional storage change [5]. A proposed quantum survey should state how it would improve upon the existing monitoring network at the same site.

The field protocol should collocate stations with selected wells, record pumping and rainfall, and repeat over a period that can capture a meaningful change. It should account for soil moisture, surface-water movement, ground elevation and instrument setup. A small pilot conducted during a stable period may return an excellent noise measurement but reveal little about seasonal aquifer storage. The contract should specify the hydrologic event or management threshold it aims to observe and a plan for independent interpretation.

The value case should be tied to a decision with money attached: whether to add a monitoring well, change pumping, approve recharge works or defer a supply project. A utility might value a non-invasive measurement where well access is limited. The cost comparison must include staff, repeated visits, land access, data integration and the wells retained for calibration. If the new survey cannot reduce the uncertainty that drives the decision, the measured gravity precision does not create a defensible investment case.

12 Test the mineral case against the drill programme

In mineral exploration, a density anomaly can sharpen a geological hypothesis about structure or a massive body, but it cannot name a commodity or establish a mineral resource. The USGS describes gravity as a method for investigating density variations over several scales and in conjunction with other geophysical methods [3]. A mineral pilot should identify the deposit model, expected density contrast, depth range, existing surveys and the drilling decision to be made. It should state which alternative interpretation would explain the same signal.

The commercial test is whether the quantum measurement changes the selection or ordering of a bounded number of targets under an equal field budget. A conventional team should submit its ranking before seeing the new result. The quantum-assisted team should then submit a revised ranking with a stated uncertainty interval. Field evidence, where legally and technically feasible, should determine whether the changed selection was warranted. The result should be reported for every selected and rejected target, including sites that could not be drilled. A supplier cannot claim credit for targets already identified in the conventional shortlist.

Ownership matters. A sensor company may own equipment and processing software but no mineral title. An explorer may own a licence but be unable to transfer the sensor data under its contract. A fund may pay for the survey yet receive no right to exploit the target. Before financing a joint venture or acquisition, counsel should map survey access, raw data ownership, derived-model rights, confidentiality and mineral title. The investor pays for the rights it receives and the field programme it can execute, rather than for a picture detached from the ground interest.

13 Test geothermal monitoring against operating data

Geothermal projects have two possible sensing questions. Exploration asks whether a prospective site warrants a costly confirmation well. Operation asks how mass moves as fluid is produced, reinjected or naturally replenished. A gravity signal may help constrain those models, but temperature, permeability, flow and chemistry require other evidence. DOE describes exploration drilling as a route to gather subsurface data and potentially confirm a geothermal resource [10]. A pre-drill gravity anomaly alone cannot establish commercial heat or sustained flow.

GFZ's Equip-G project plans campaigns using an absolute quantum gravimeter in a geothermal field, alongside other instruments, to study short-term mass redistribution [6]. The project is a relevant research pathway rather

than proof that a commercial deployment has achieved a particular operating saving. A developer considering a pilot should choose a site with known production and injection records, existing gravity or geodetic baselines, and an operating event with a predicted measurable response. The survey should include background reference stations and an independent model of other mass changes.

The economic decision may concern a reinjection strategy, the timing of a further well or the size of a reservoir-management programme. The sensor's output must be timely and sufficiently interpretable to affect that decision. A financing model should hold project cash flow separate from the value of the monitoring contract. If a vendor claims a share of future electricity revenue, the buyer should test how much of that revenue the measurement actually caused, and whether the claim survives a change of instrument supplier.

14 Test storage monitoring against the approved plan

Geological carbon storage is a monitoring obligation as well as an information problem. NETL describes subsurface monitoring as a way to track the carbon dioxide plume, physical-property changes and potential migration pathways [7]. It lists high-precision gravity among several tools and notes the difficulty of monitoring large areas. PNNL reviews time-lapse gravity as a method for observing subsurface mass redistribution at storage sites [9]. These sources support a technical role for gravity. They do not authorize a quantum device to replace a site's approved monitoring plan.

A storage pilot should begin with the site-specific risk assessment, injection history and monitoring commitments. The sponsor should ask whether added gravity observations could reduce uncertainty about plume extent, pressure-related changes or mass balance. The sensor should be compared with seismic, wells, pressure, tracers and other methods that address the same risk. Interpretation must distinguish injected fluid effects from brine displacement, groundwater change and geological or operational noise. Regulators and technical reviewers determine the acceptance requirements for the actual site.

The value case may include a more efficient survey cadence or better targeting of confirmatory monitoring, but those benefits need site evidence. The buyer should avoid placing a hypothetical reduction in long-term liability into current project value without an approved method and measured performance. A contract can support a pilot as additional evidence, preserve all raw data for independent review and defer any recurring service commitment until the result has been evaluated against the site's monitoring plan.

15 Use a fictional case to expose the economics

Consider a fictional operator with six candidate subsurface test sites. Its current plan includes two intrusive confirmation tests that might prove unproductive. The author assumes that each avoidable unsuccessful test would cost USD 1.5 million. A proposed sensor survey, interpretation and independent validation cost USD 0.9 million in total. The author assumes a 70% chance that the survey correctly prevents each of the two avoidable tests. The author also assumes a 10% programme-level chance that relying on the survey causes the operator to miss one viable opportunity with a conditional value of USD 4 million. These assumptions are invented solely to expose the decision arithmetic.

The expected gross avoided cost is two times 70% times USD 1.5 million, or USD 2.1 million. Subtracting the USD 0.9 million survey cost and the 10% times USD 4 million missed-opportunity allowance leaves USD 0.8 million of illustrative incremental value. That figure is not a sensor price, a probability estimated from field trials or a return available to a real operator. It excludes financing delay, correlation among site outcomes, replacement surveys, ownership rights and any contractual penalty. Those items would need to be modelled for an actual transaction.

The simple case is useful because it identifies the evidence the buyer needs. A blind field trial must support the avoided-test rate for the specific geology and threshold. A review of negative and ambiguous results must constrain the missed-opportunity risk. Procurement must establish the full installed survey cost. The site owner must document the monetary consequence of a wrong decision. If these inputs are unavailable, the buyer can still fund a bounded pilot; it cannot defend a scaled valuation using this educational calculation.

Table 2. Fictional six-site decision case under author assumptions

Case input or output	Author assumption or calculation	Basis in this example
Avoidable unsuccessful tests	2	Hypothetical programme choice
Cost of each unsuccessful test	1.50	Assumed all-in intrusive test cost
Chance of correctly avoiding each test	70%	Assumed, not a field result
Sensor survey and validation cost	0.90	Assumed complete pilot cost
Missed-viable-site event chance	10%	Assumed programme-level risk
Conditional loss if missed	4.00	Assumed lost option value
Gross expected avoided cost	2.10	$2 \times 0.70 \times 1.50$
Illustrative incremental value	0.80	$2.10 - 0.90 - 0.40$


USD millions unless stated otherwise; no line is a measured commercial result or observed sensor performance.

16 Stress the false negative before paying for scale

A sensor may appear valuable if the model counts only avoided unsuccessful tests. The cost of missing a viable site can reverse that result. In the fictional case, the break-even probability of correctly avoiding each of the two tests is about 43.3% when the survey costs USD 0.9 million and the missed-opportunity allowance remains USD 0.4 million. At a 70% avoidable-test rate, raising the programme-level miss probability from 10% to 20% reduces illustrative incremental value from USD 0.8 million to USD 0.4 million. A larger conditional loss or higher survey cost can eliminate the gain.

These sensitivities should be contractual questions, not curve-fitting after a demonstration. What counts as a missed viable site? Who adjudicates a target that cannot be drilled? How long after the survey may ground truth arrive? Does the supplier's fee fall when the survey does not change the decision? A deployment business is worth more when its performance remains useful under realistic failure definitions and when it can bear an appropriate share of its controllable risks.

An investor can set a maximum first-stage commitment as the cost of obtaining reliable evidence about those rates. The pilot need not be expected to make a profit on its own; it can be a rational expense if it prevents a much larger misallocation. That claim requires a plausible route to repeatable customers and a clear decision to stop after a poor result. The model should keep sensor-company enterprise value separate from a customer's project value so that a single assumed saving does not get capitalised twice.

Figure 2. Fictional incremental value under avoided-test and missed-site assumptions

All rates and USD values are author assumptions; the chart is a decision sensitivity, not a measured sensor-performance curve.

17 Compare business models on cash and rights

A hardware sale puts acquisition and maintenance risk with the customer. A survey-service contract keeps instrument ownership and field operations with the supplier. A recurring data product may combine deployment with interpretation and archiving. A joint venture can link the provider's compensation to a field programme. Each can be sensible under different evidence and capital needs. The investor should compare gross margin after travel, skilled operators, calibration, spare equipment, insurance, data processing, interpretation and customer support, not only the instrument's headline selling price.

The revenue contract must say who owns raw gravity observations, corrected datasets, site models, software updates and derived target maps. A mine or utility may need unrestricted access to its site data for later financing or regulation. A vendor may need reusable algorithms and non-sensitive learning. These interests can be reconciled through a clear licence and audit trail. A financing party should inspect customer termination rights, indemnities, export controls, public-sector procurement restrictions and the ability to assign contracts in an acquisition.

The scale thesis also requires fleet arithmetic. If one instrument requires several days of mobilisation for a short survey, utilisation may be low even when the measurement is useful. A company can improve economics through repeat routes, local service partners and standardised processing, but those are operational hypotheses until demonstrated. The diligence model should state number of working instruments, available crews, survey days, rework rate, payable deliverables and cash-collection terms. Revenue promised for a future fleet should not be treated as contracted business.

18 Allocate error and follow-on risk

The supplier controls instrument condition, calibration, data capture, processing and compliance with the agreed survey protocol. The site owner controls access, disclosed site records and its subsequent intrusive tests. Independent interpreters may control the geological model. Weather, unrecorded infrastructure, hydrologic changes and deeper geology can defeat all parties' expectations. A sensible contract maps obligations to the actor that can manage them and defines the limits of each warranty.

A sensor vendor can warrant that it delivered a specified dataset and reproducible uncertainty estimate. It should not promise that a particular aquifer, orebody or reservoir exists merely because the instrument registered an anomaly. The customer should warrant that it supplied the site information it agreed to provide, not that every hidden feature is known. Where performance-linked consideration is proposed, the endpoint must be independent of the vendor's preferred interpretation and achievable within a defined observation period. Disputes over inconclusive results need a named expert and a bounded resolution process.

Capital providers should see a loss tree. It should show pilot cost if the instrument fails, additional cost if interpretation is ambiguous, the consequence of a false positive and the consequence of a false negative. It should specify who funds a repeat survey and whether the operator may use another method without losing rights to the initial data. This makes the investment contract resilient to the most common outcome of early subsurface work: information that narrows uncertainty without resolving every question.

19 Diligence the sensor company and the project separately

The sensor company needs technical diligence on atom source, lasers, vacuum system, environmental tolerance, calibration, quality control, software and field uptime. A buyer should inspect test logs across the full fleet, not a selected laboratory instrument. It should ask how the unit is transported, levelled and maintained; how failure is diagnosed; and whether replacement parts or specialised staff create a single point of failure. Benchmarks should report complete survey throughput and usable data yield. The 2022 tunnel paper gives an experimental reference, not a warranty for a future commercial product [1].

The customer's underground project needs different diligence: title, access, geological setting, existing geophysics, well or excavation plans, environmental constraints, operating records and the value of the next decision. The instrument company can be excellent while a particular site offers little measurable density contrast. The site can be attractive while the sensor company lacks the capacity to deliver a reproducible survey. A financing model that merges these two risks into one optimistic probability hides the reason a deal might fail.

An independent geophysicist should review the forward model and alternatives. A water, mining, geothermal or storage specialist should review the domain-specific truth test. Counsel should review data and site rights. A commercial reviewer should confirm that the buyer of the measurement has budget authority and will pay for a result with the agreed uncertainty. The final diligence file should link every major valuation input to an observed contract, test result, source or labelled author scenario. Missing evidence remains missing; it does not become zero or a market average by convenience.

Table 3. Investment committee tests before a commercial sensor commitment

Gate	Evidence to request	Decision if unresolved
Physical target	Forward model and competing density explanations	Redesign or reject the site test
Field performance	Blind survey and full station-level error budget	Limit funding to validation
Decision benefit	Comparator, truth test and cost of wrong action	Do not price a performance premium
Rights and operations	Data rights, site access, fleet and support plan	Withhold scaled service commitment
Customer demand	Signed scope, budget owner and payment terms	Keep revenue in a scenario

Author-developed governance checklist; thresholds must be set for the actual customer and site.


20 Move through bounded evidence and payment gates

The first gate is an application screen. The buyer and an independent geophysicist define the physical target, expected signal, site access, conventional comparator and value of the decision. If the target cannot produce a discriminating signal under plausible assumptions, the programme stops before equipment is deployed. The second gate is a blind field pilot. Data are frozen and independently processed, with all stations and adverse results retained. Payment at this stage covers a defined service and data package.

The third gate is independent truth and decision review. The committee compares the sensor-assisted choice with the conventional plan and records whether the sensor changed the action correctly. A shallow target may be checked quickly; groundwater, geothermal and storage monitoring may require repeated observations and a longer event window. The fourth gate is a scale contract, which should be signed only after the unit economics, data rights, operational capacity and customer budget are evidenced. A good pilot can warrant further research without warranting an acquisition price or a multi-year exclusivity commitment.

The roadmap should contain explicit stop conditions, a maximum cash commitment and a named owner for each piece of evidence. A deposit or reservoir project should retain its own technical approvals. If a regulator or operator requires a separate monitoring method, the sensor contract must accommodate it. A staged structure protects both parties: the supplier can be paid for real deliverables, and the customer can reserve larger payments for independently demonstrated decision value. The transaction remains capable of learning from failure.

At each gate, the committee should update the same decision ledger rather than begin a new sales narrative. The ledger records the prior probability range, the measurement obtained, the alternative explanations, the independent check, the action actually taken and the cash spent. A rejected target belongs in the record alongside a successful one. This discipline matters because a sensor programme may appear more accurate when only the best sites reach a final truth test. A customer considering a second geography should request the full first-geography ledger, including inaccessible sites, failed stations and inconclusive interpretations. The next contract can then specify which parts of the workflow transfer and which require fresh validation.

Figure 3. Author-developed four-gate commercial test and payment route

Stop, redesign or repeat within an approved cap after every gate

The route is a governance sequence; it contains no forecast of technical success or approval time.

21 Conclusion

Quantum gravity sensing has crossed a meaningful experimental threshold in an outdoor tunnel survey [1]. Its commercial value remains a site-specific proposition. The investor needs a physical model of the target, a credible conventional comparator, a blind field test, independent truth and a decision that changes enough to justify the complete deployment cost. Water, mineral, geothermal and storage projects require different validation periods and different economic endpoints. Instrument sensitivity should be translated into those endpoints only through evidence.

The fictional six-site case shows why both kinds of error matter. Avoiding unsuccessful tests can create value; missing a viable site can consume it. The case's USD 0.8 million illustrative incremental value rests entirely on author assumptions and does not describe a real sensor, project or company. The break-even test identifies the field evidence that would need to be collected before a price or recurring fee could be defended. No resource or storage asset should be booked merely because a gravity image has become more detailed.

A workable mandate begins with the next decision and a bounded pilot. It pays for delivered measurement and reproducible interpretation, reserves larger payments for independently confirmed decision benefit, and retains the right to stop when evidence is adverse. That structure gives a sensor supplier a route to commercial proof and gives the capital provider a clear limit on the risk taken before underground truth is available.

Frequently Asked Questions

What has quantum gravity sensing demonstrated outdoors?

A published cold-atom gravity-gradiometer field survey detected a known buried tunnel under the conditions described in that experiment. It does not establish commercial performance for every underground target.

Can a gravity anomaly identify a specific mineral or aquifer?

No. Gravity responds to mass distribution and density contrast. Geological context, other measurements and independent field tests are needed to identify the cause and commercial relevance of an anomaly.

Which of the four applications has the strongest immediate investment case?

The strongest case is the particular site where a measurement can be independently validated and changes a costly decision. The paper does not rank water, minerals, geothermal and storage as universal markets.

How should a sensor be compared with conventional methods?

Compare complete field workflows against the same site, decision, access, time window and truth test. Include survey, processing, interpretation, maintenance and the costs of false results.

What does the fictional USD 0.8 million figure mean?

It is an author-assumed illustration of avoided test cost after survey cost and missed-site risk. It is not an observed saving, vendor price or investment return.

Does the published tunnel result justify using a sensor for deep geothermal or carbon storage?

It justifies investigating a site-specific application. Deep geothermal and storage proposals need their own forward model, field pilot, comparator and independent operating evidence.

Who should own the survey data?

The parties should define raw observations, processed data, site interpretation, software and reuse rights in the contract. The buyer must retain the evidence needed for operation, financing and any regulatory review.

When should an investor fund a sensor fleet rather than a pilot?

A fleet commitment requires repeatable field performance, customer contracts, support capacity, utilisation and unit economics. A successful bounded pilot alone does not establish those conditions.

References

1. Stray, B. et al. Quantum sensing for gravity cartography. *Nature* 602, 590-594 (2022). Accessed 5 October 2026. <https://www.nature.com/articles/s41586-021-04315-3>
 2. University of Birmingham. Quantum Technology Sensing, National Buried Infrastructure Facility. Accessed 5 October 2026. <https://www.birmingham.ac.uk/research/centres-institutes/uk-collaboratorium-for-research-on-infrastructure-and-cities/national-buried-infrastructure-facility-nbif/nbif-research-capabilities/quantum-technology-sensing>
 3. US Geological Survey. Tools and Techniques: Gravitational Method. Accessed 5 October 2026. <https://www.usgs.gov/publications/1112-tools-and-techniques-gravitational-method>
 4. Zohdy, A. A. R., Eaton, G. P. and Mabey, D. R. Application of Surface Geophysics to Ground-Water Investigations. US Geological Survey (1974). <https://www.usgs.gov/publications/application-surface-geophysics-ground-water-investigations>
 5. US Geological Survey. Arizona Water Science Center: Hydrologic Gravity Monitoring. Accessed 5 October 2026. <https://www.usgs.gov/centers/arizona-water-science-center/science/azwsc-capabilities-hydrologic-gravity-monitoring>
 6. GFZ German Research Centre for Geosciences. EQUIP-G: Absolute Quantum Gravimetry for Geothermal Resources Monitoring. Accessed 5 October 2026. <https://www.gfz.de/en/section/geophysical-imaging/projects/equip-g-absolute-quantum-gravimetry-for-geothermal-resources-monitoring>
 7. US Department of Energy, National Energy Technology Laboratory. Subsurface Monitoring. Accessed 5 October 2026. <https://www.netl.doe.gov/node/5873>
 8. US Geological Survey. Historical Development of the Gravity Method in Exploration. Accessed 5 October 2026. <https://www.usgs.gov/publications/historical-development-gravity-method-exploration>
 9. Pacific Northwest National Laboratory. Monitoring Carbon Storage Sites with Time-Lapse Gravity Surveys. Accessed 5 October 2026. <https://www.pnnl.gov/publications/monitoring-carbon-storage-sites-time-lapse-gravity-surveys>
 10. US Department of Energy. Pilots and Demonstrations: Next-Generation Field Tests and Drilling for New Resources. Accessed 5 October 2026. <https://www.energy.gov/hgeo/geothermal/pilots-and-demonstrations-next-generation-field-tests-and-drilling-new-resources>
 11. US National Institute of Standards and Technology. Turning Atoms Into Waves to Measure Gravity and Acceleration. Accessed 5 October 2026. <https://www.nist.gov/quantum-information-science/quantum-sensing-explained/turning-atoms-waves-measure-gravity-and>
- Sources reviewed 5 October 2026. The six-site example, probabilities, costs, values, figures, tables and commercial gates identified as fictional or author-developed are educational decision tools. They are not observations of sensor performance, transaction prices, project returns or geological resources. This paper provides no geological, engineering, legal, tax, accounting, regulatory or investment opinion on a particular property, vendor or transaction.

About the Author

Authored by Chennakeshav Adya
Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.