

Beyond Compute Futures: The Economics of Computational Risk and Customized Derivatives

AUTHORS: Peter C. Earle, Ph.D

SUMMARY: This paper explores the economic implications of computational capacity as a critical production input, highlighting the increasing uncertainty firms face regarding its price, availability, and technological characteristics. The author argues that while standardized compute futures can address some benchmarkable risks, substantial computational risks remain too heterogeneous for effective exchange-traded hedging. By drawing on the evolution of various derivative markets, the paper develops a framework for understanding when computational risk favors customized bilateral contracts.

METHOD: The study draws on the historical development of commodity, electricity, freight, and interest-rate derivatives to establish a framework for understanding computational risk. It examines the potential structures of compute swaps and the benchmarks and market infrastructure required to support them.

KEY FINDINGS:

- Standardized compute futures can hedge benchmarkable price exposures but fail to address the heterogeneous nature of computational risks.
- The division between standardized futures and customized derivatives is determined by the underlying heterogeneity of computational risk.
- Compute swaps are necessary for managing specific risks related to hardware, geography, workload, and contractual structure.
- Financially settled H100 and B200 rental index futures and GPU compute futures are examples of standardized compute futures.

IMPLICATION FOR TRADING: The paper suggests that traders and firms should consider the limitations of standardized compute futures and explore customized bilateral contracts for managing specific computational risks. This approach can help in reducing basis risk and ensuring more tailored risk management strategies.

Beyond Compute Futures: The Economics of Computational Risk and Customized Derivatives

Peter C. Earle, Ph.D
Senior Director of Research
American Institute for Economic Research
pete.earle@aier.org

19 September 2026

WORKING PAPER - DRAFT

Abstract:

Computational capacity has become an increasingly important production input, exposing firms to uncertainty over prices, availability, technological change, and infrastructure utilization. This paper argues that while standardized compute futures can hedge increasingly benchmarkable price exposures, substantial computational risks remain too heterogeneous for effective exchange-traded hedging. Drawing on the evolution of commodity, electricity, freight, and interest-rate derivatives, it develops a framework for understanding when computational risk favors customized bilateral contracts. The paper identifies several potential compute swap structures and examines the benchmarks and market infrastructure required to support them. More broadly, it argues that the division between standardized futures and customized derivatives will be determined by the underlying heterogeneity of computational risk.

JEL Classifications: D24, D47, D82, G13, G14, L86, O33

Keywords: computational risk; compute swaps; computational capacity; over-the-counter derivatives; computational finance; artificial intelligence; cloud computing; institutional economics; risk management

1. Introduction

Computational capacity has become one of the defining productive inputs of the twenty-first century. Once regarded primarily as an operational expense associated with information technology, compute now underpins artificial intelligence, cloud services, engineering simulation, pharmaceutical discovery, scientific research, financial modeling, and a growing array of data-intensive commercial activities. Increasingly, firms compete not only for labor, capital, and energy, but also for access to specialized computational resources capable of executing increasingly sophisticated workloads. For many organizations, access to computational capacity has become as economically consequential as access to more traditional factors of production.

This transformation has been accompanied by an equally important change in the nature of economic risk. Organizations committing billions of dollars to artificial intelligence, cloud infrastructure, and computational research must make long-term investment decisions despite substantial uncertainty regarding the future price, availability, and technological characteristics of the computational resources upon which those investments depend. These uncertainties increasingly resemble the forms of commercial risk that have historically given rise to derivative markets.

Derivative markets tend to appear when commercially important uncertainty becomes worth transferring. Agricultural forwards and futures grew around price and inventory risk; organized futures markets lowered the costs of trading standardized exposures (Working 1949; Telser 1981). Corporate derivatives later expanded into risks that were less naturally standardized, particularly interest rates and currencies, where hedging could protect investment and financing plans from unwanted volatility (Stulz 1984; Froot, Scharfstein, and Stein 1993). Compute now presents a related problem, but with an unusually heterogeneous underlying resource.

The relevant uncertainty is not confined to the hourly rental price of a GPU. Semiconductor supply, electricity costs, data center utilization, reservation terms, regional congestion, and successive hardware generations can all alter the economics of a computational project. A firm may know how much work it expects to perform while remaining uncertain about the price, location, architecture, or availability of the capacity required to perform it. Cloud-computing research has already treated forward pricing as a mechanism for managing stochastic computational demand, and earlier market design work showed that computational resources can support explicit allocation and pricing institutions (Bapna et al. 2008, 2011; Du, Das, and Ramesh 2013).

That combination of economic importance and heterogeneity is the problem examined here. Compute is becoming measurable enough for some risks to be standardized, but not homogeneous enough for a single contract to represent every commercially relevant exposure. Hardware architecture, memory, networking, software compatibility, latency, location, workload, and contract terms can matter independently. The resulting basis risk is not incidental. It helps determine which risks can migrate to an exchange and which remain better suited to bilateral contracting.

This question is no longer hypothetical. CME Group and Silicon Data have announced financially settled H100 and B200 rental index futures, while Nodal Exchange and Compute Desk have announced GPU compute futures tied to published rental price indexes (CME Group 2026a, 2026b; Nodal Exchange 2026). Those contracts support the argument in Earle (2026) that sufficiently benchmarkable compute can support cash-settled futures. They also expose the residual question: what happens to computational risks that a standardized rental price contract does not hedge?

The proposition developed here is that the boundary between futures and swaps will be set by standardization and basis risk. As computational expenditures become larger and more forward-committed, demand for risk transfer mechanisms should increase. Where a credible benchmark represents the underlying exposure closely enough, standardized futures can concentrate liquidity and lower trading costs. Where material exposures remain specific to hardware, geography, workload, availability, utilization, technology generation, or contractual structure, customization can be worth its higher transaction and credit costs. Compute swaps should therefore be most useful where the exposure is economically significant but only imperfectly represented by an exchange benchmark.

That is a narrower claim than the one made in Earle (2026). The earlier paper considered the broader financialization of computational capacity. This paper concentrates on the risks left over after standardization has done what it can. Futures and swaps are therefore treated as complementary institutions, not stages in a mandatory sequence. The choice between them follows from the characteristics of the exposure and the costs of contracting, a point consistent with transaction cost approaches to institutional form (Coase 1937; Williamson 1979).

2. Why Computational Capacity Requires Derivative Markets

The case for computational derivatives begins with the scale and duration of the commitments being made. Firms increasingly undertake AI training, scientific computing, engineering, cloud migration, and infrastructure projects whose computational requirements extend well beyond the date on which capital is committed. The commercial

problem is familiar: an investment decision is made today while an important input cost remains uncertain. Corporate hedging theory provides a straightforward reason to care. If volatile input costs disrupt investment or financing plans, transferring some of that risk can have value even when shareholders themselves are diversified (Stulz 1984; Froot, Scharfstein, and Stein 1993).

Compute complicates the familiar commodity analogy because it is not simply a fungible input with a volatile spot price. Pharmaceutical simulation, quantitative finance, rendering, frontier-model training, and industrial digital twins can require different hardware, memory, networking, latency, security, and software environments. The economically relevant unit is therefore partly technological and partly contractual. That is why computational risk cannot be reduced to one price series without losing information that matters to the user.

Some of that complexity can nevertheless be standardized. Published rental indexes can support a hedge against broad GPU price movements, just as commodity benchmarks permit firms to hedge exposures that are similar, though rarely identical, to the deliverable or settlement asset. Commodity theory has long emphasized the relation among spot prices, inventories, forward prices, and risk bearing (Brennan 1958; Black 1976; Fama and French 1987; Hirshleifer 1988). Compute adds an unusually fast-moving technological dimension to the familiar basis problem.

The practical issue is how much mismatch a hedger can tolerate. A national H100 rental index may be useful to a buyer whose costs move closely with that benchmark. It may be a poor hedge for a firm requiring capacity in one region, under a particular reservation structure, with specific networking or software characteristics. Likewise, a data center operator may care less about the rental price alone than about the spread between revenue, electricity costs, financing costs, and utilization. The same underlying industry can therefore generate several distinct hedgeable exposures.

This is where bilateral contracts retain an economic role. Standardization creates liquidity by suppressing contractual differences. Customization sacrifices some of that liquidity in order to match the exposure more closely. The tradeoff resembles the one observed in other derivative markets in which liquid benchmark instruments coexist with basis swaps, spread contracts, and other over-the-counter agreements. There is no reason to expect compute to resolve that tradeoff differently simply because the underlying resource is digital.

3. The Boundary Between Futures and Swaps

Futures and swaps solve related but different coordination problems. Organized futures markets work best when many participants are willing to trade the same contract. Standard

quantities, settlement rules, maturities, and reference assets concentrate orders and make prices observable. Telser (1981) emphasized the institutional economies of organized futures markets, while Black (1976) provided the canonical pricing framework for commodity contracts. Those advantages are substantial, but they depend on a contract that is close enough to the exposures participants actually want to hedge.

A swap gives up some of those economies in exchange for fit. The parties can choose the index, quantity profile, location, observation period, technology specification, settlement rule, and other terms. That flexibility is valuable only when the reduction in basis risk exceeds the additional costs of negotiation, collateral, credit management, valuation, and reduced liquidity. The relevant comparison is therefore not customization versus standardization in the abstract. It is the marginal value of a better hedge against the marginal cost of obtaining it.

Compute makes that comparison unusually visible. A standardized H100 rental future can hedge a common price factor. It cannot, by itself, guarantee capacity, hedge the spread between two regions, protect against a migration from one hardware generation to another, or stabilize the operating margin of a data center owner. Those exposures may be correlated with the benchmark and therefore partly hedgeable through futures. The residual can still be large enough to matter.

The market design implication follows from that residual rather than from any presumed chronology. Exchange trading should expand where performance is comparable, prices are transparent, and basis risk is tolerable. Bilateral contracts should survive where users value terms that cannot be standardized without destroying hedge effectiveness. Improvements in benchmarks can move the boundary in either direction. A risk that requires a bespoke contract today may become standardized later; a new hardware architecture or commercial practice can create new basis risks just as quickly.

This also places a limit on the case for compute swaps. Heterogeneity alone does not create a viable derivative. The exposure must be measurable, settlement must be verifiable, and there must be a counterparty willing to take the other side at a price that compensates for risk and contracting costs. Grossman and Stiglitz (1980) remind us that information is costly, and computational markets may make some relevant characteristics expensive to observe. Where measurement is poor or disputes over performance are likely, customization may produce an untradeable contract rather than a superior hedge.

Nor must every computational risk be financialized. Long-term cloud contracts, reserved instances, vertical integration, ownership of hardware, geographic diversification, or simply carrying residual risk can substitute for derivatives. Coase (1937) and Williamson (1979)

provide the broader institutional point: market participants choose among contractual and organizational arrangements rather than derivatives in isolation. A compute swap is attractive only when it reallocates risk more cheaply than those alternatives.

The claim of this paper is consequently conditional. Compute swaps have a plausible domain where exposures are material, measurable, and poorly represented by standardized contracts, and where the value of a closer hedge exceeds the costs of bilateral contracting. The taxonomy that follows identifies the principal risks for which those conditions might hold.

4. A Taxonomy of Computational Risk

Computational risk is not one uncertainty. A firm can be right about the future rental price of compute and still be wrong about whether the needed hardware will be available, where it will be available, whether a new architecture will displace it, or whether utilization will cover fixed costs. Separating these exposures is useful because they need not be hedged together.

The categories below are therefore economic exposures rather than predictions about contract names. Some may eventually be absorbed into standardized futures or options. Others may remain bilateral, be handled through ordinary commercial contracts, or prove too difficult to measure for a derivatives market at all.

4.1 Price Risk

The most obvious exposure is uncertainty surrounding the future cost of computational resources. Organizations committing to long-term artificial intelligence training, scientific research, cloud migration, or engineering projects frequently do so without knowing the future price of the computational services those projects will require. Future compute prices depend upon semiconductor availability, electricity costs, capital expenditures, cloud provider pricing policies, and shifts in demand across industries. Rising compute prices may substantially increase project costs long after investment decisions have been made.

Price risk resembles the commodity price risks historically transferred through energy and raw material derivatives. A compute swap allowing counterparties to exchange floating compute prices for predetermined fixed prices would permit organizations to stabilize budgets while allowing providers or financial counterparties to assume the associated market risk.

4.2 Capacity Availability Risk

Price is only one dimension of computational uncertainty. During periods of elevated demand, specialized hardware may simply become unavailable regardless of price. Recent shortages of advanced GPUs demonstrated that organizations were frequently unable to obtain computational resources necessary to complete planned projects even when willing to pay prevailing market rates.

Availability risk resembles the capacity risks observed in electricity markets and transportation networks, where access to infrastructure can become more economically significant than the nominal price of the underlying service. Customized agreements guaranteeing future computational availability could therefore hedge risks associated with delayed projects, postponed product launches, or interrupted research programs.

4.3 Technology and Obsolescence Risk

Computational infrastructure evolves unusually rapidly. Successive processor generations frequently deliver substantial improvements in performance, energy efficiency, memory architecture, and networking capability. Organizations purchasing or leasing computational resources therefore face uncertainty not only regarding future prices but also regarding the relative value of existing assets as technology advances.

This exposure differs fundamentally from ordinary depreciation. Rapid technological innovation may reduce the economic value of computational assets long before their physical lives have ended. Compute swaps indexed to benchmark performance measures or hardware generations could permit firms to hedge the financial consequences of technological transition while retaining operational flexibility.

4.4 Geographic and Basis Risk

Computational services are increasingly provided through geographically dispersed data centers exhibiting different pricing structures, energy costs, regulatory environments, and infrastructure constraints. Consequently, organizations operating across multiple jurisdictions often face regional pricing differentials analogous to the basis risks observed in energy, freight, and agricultural markets.

Basis risk arises whenever the exposure being hedged differs from the benchmark underlying a financial contract. Customized swaps indexed to regional computational benchmarks could reduce these mismatches by allowing counterparties to exchange localized pricing exposures without requiring universal standardization.

4.5 Utilization and Operating Margin Risk

For providers of computational infrastructure, profitability depends not only upon compute prices but also upon utilization rates and operating costs. Idle servers generate little

revenue while continuing to incur capital, maintenance, and energy expenses. Conversely, unexpectedly high utilization may create opportunities for significant operating profits.

Infrastructure providers therefore confront risks analogous to those faced by airlines, hotels, electric utilities, and shipping companies, whose profitability depends upon both pricing and capacity utilization. Financial contracts linked to utilization rates or operating margins could permit providers to stabilize cash flows while transferring demand fluctuations to counterparties better positioned to absorb them.

4.6 Portfolio Risk

Few organizations face only one of the foregoing risks. Most simultaneously confront changing compute prices, uncertain capacity, evolving hardware, regional pricing differences, and fluctuating utilization. These exposures rarely move together and often affect organizations differently depending upon their business models, technological strategies, and contractual arrangements.

5. Compute Swap Architectures

The taxonomy can be translated into several plausible contract structures. These are not forecasts of a definitive product menu. They illustrate how a counterparty might isolate a computational exposure without taking delivery of the underlying service.

The analogy is with the proliferation of specialized contracts in established derivatives markets. Interest rate markets developed basis, inflation, cross-currency, and amortizing swaps; electricity and freight markets developed contracts around locations, routes, capacity, and spreads. Compute could generate similar specialization if the underlying risks become large enough to justify the contracting costs (Hull 2022).

5.1 Fixed-for-Floating Compute Price Swaps

The simplest compute derivative would closely resemble a conventional commodity or interest rate swap. One counterparty would agree to pay a predetermined fixed computational price while receiving payments based upon prevailing market prices over a specified period.

Such contracts would primarily benefit organizations seeking budget certainty. Large artificial intelligence developers, pharmaceutical firms, engineering companies, and research institutions frequently commit to projects whose computational requirements extend months or years into the future. Locking in future compute costs would stabilize project economics despite changing market conditions.

Cloud providers or specialized financial intermediaries could assume the opposite exposure, profiting when realized market prices differ favorably from the negotiated fixed price.

Unlike exchange-traded futures, settlement would almost certainly remain cash-based rather than involving physical delivery of computational services. Operational procurement would continue independently of financial settlement.

5.1.1 A Simple Hedging Example

Consider an AI developer that expects to consume 10 million benchmark compute units during the next year. Its operational procurement remains with its cloud or GPU provider, but management wants to stabilize the financial cost of that consumption. Suppose the firm enters a fixed-for-floating compute swap at a fixed benchmark price K of \$2.00 per unit for a notional quantity Q of 10 million units. At settlement, the swap payment to the firm is $Q \times (P - K)$, where P is the observed benchmark price. If the benchmark rises to \$2.50, the swap pays the firm \$5 million: 10 million $\times$ (\$2.50 - \$2.00). That gain offsets, approximately, the higher cost of acquiring compute in the operating market. If the benchmark falls to \$1.50, the firm pays \$5 million under the swap but benefits from lower procurement costs. The contract therefore converts a floating benchmark exposure into an approximately fixed financial cost without requiring physical delivery of GPUs or cloud services.

The qualification is important. The hedge is exact only if the firm's actual procurement exposure moves one-for-one with the settlement benchmark. If the firm requires H100-equivalent capacity in Northern Virginia with particular latency, networking, reservation, or software characteristics while the swap settles against a broader national rental index, the difference remains as basis risk. That residual is not a defect in the argument for computational derivatives; it is the reason customized swaps have a role alongside standardized futures. A bilateral contract can alter the index, location, hardware specification, observation period, quantity profile, or other terms to reduce the mismatch between the financial hedge and the firm's actual computational exposure.

5.2 Basis Swaps

Many organizations face not absolute price risk but differences between computational benchmarks.

A multinational firm may purchase compute capacity in several geographic regions. A cloud provider may price services using one benchmark while customers purchase capacity indexed to another. Future market development may produce independent

benchmark providers whose indices differ systematically according to hardware generation, location, or workload characteristics.

Basis swaps would permit counterparties to exchange one computational benchmark for another without hedging the absolute level of compute prices themselves. Such instruments already play an important role in electricity, petroleum, freight, and interest rate markets whenever pricing differs across otherwise related assets.

Because computational markets remain geographically and technologically fragmented, basis swaps are likely to become an important component of future computational finance.

5.3 Technology Transition Swaps

Perhaps the most distinctive feature of computational markets is rapid technological change.

Organizations investing heavily in today's hardware face uncertainty regarding the pace at which future processor generations may reduce the relative economic value of existing infrastructure. Conversely, cloud providers introducing new architectures face uncertainty regarding customer adoption and the resulting utilization of legacy systems.

Technology transition swaps could allow counterparties to exchange financial exposures associated with changing hardware generations, effectively separating operational decisions regarding technology adoption from the financial consequences of technological obsolescence.

Such contracts have few direct analogues in traditional commodity markets, reflecting the unusually rapid innovation characteristic of computational infrastructure.

5.4 Capacity Availability Swaps

During recent periods of exceptionally strong demand, computational capacity became difficult to obtain regardless of price. Organizations experienced delays not because they were unwilling to pay prevailing market rates but because physical capacity simply did not exist.

Capacity availability swaps would compensate purchasers whenever reserved computational resources became unavailable according to specified contractual criteria.

Unlike conventional insurance, however, these contracts would remain bilateral financial agreements negotiated between commercial counterparties rather than contingent indemnification arrangements.

The distinction mirrors capacity contracts already observed in electricity markets where infrastructure availability often proves more valuable than the underlying commodity itself.

5.5 Operating Margin Swaps

Cloud providers and data center operators derive profitability from the interaction of compute prices, electricity costs, utilization rates, financing costs, and capital expenditures.

Rather than hedging any single variable individually, operating margin swaps would allow providers to stabilize overall profitability.

These contracts resemble refinery margin ("crack") spreads, spark spreads in electricity markets, and other industrial margin derivatives already used extensively in commodity markets.

Because computational infrastructure increasingly resembles capital-intensive industrial production, similar financial instruments may emerge naturally as markets mature.

5.6 Portfolio Swaps

Large organizations rarely confront only one computational exposure.

Instead, they simultaneously face changing compute prices, uncertain hardware availability, regional pricing differences, evolving processor generations, and fluctuating utilization.

Portfolio swaps would bundle multiple computational risks into a single customized agreement reflecting the particular operating characteristics of an individual organization.

These instruments would most closely resemble highly customized over-the-counter transactions negotiated between sophisticated institutional counterparties rather than standardized financial products.

Discussion

Taken together, these examples reinforce a broader institutional conclusion.

Computational finance is unlikely to develop around a single universal derivative contract.

Rather, market participants are likely to construct families of related swap agreements, each designed to isolate particular dimensions of computational risk.

Financial innovation has historically proceeded in precisely this manner. Markets evolve incrementally as participants identify economically important risks and develop contractual mechanisms capable of reallocating them more efficiently. Computational capacity appears increasingly likely to follow that same institutional trajectory.

6. Market Infrastructure and Benchmark Construction

The existence of economically meaningful risk and an appropriate contractual framework does not, by itself, produce a functioning derivatives market. Successful financial markets depend upon supporting institutions capable of measuring underlying exposures, establishing transparent benchmarks, facilitating settlement, managing counterparty risk, and generating sufficient confidence among participants to encourage trading. Throughout financial history, these institutions have frequently developed alongside derivative markets rather than preceding them. Compute swaps would likely follow the same evolutionary pattern.

Importantly, many of the institutional prerequisites for computational finance are already emerging. Cloud providers increasingly publish transparent pricing schedules for standardized computing services, benchmark organizations continue refining objective measures of computational performance, and independent firms collect detailed information on hardware characteristics, utilization, energy consumption, and market pricing. Although these developments have been driven primarily by operational rather than financial considerations, they collectively provide much of the informational infrastructure upon which derivative markets traditionally depend. The challenge is therefore less one of inventing entirely new institutions than adapting existing financial market architecture to a new class of economic exposure.

6.1 Benchmark Construction

Every derivative needs an observable reference variable. Agricultural futures use grades and delivery specifications; interest rate swaps use published rates; electricity and freight contracts use hubs, nodes, and route indexes. The informational function is important as well as the mechanical one. Prices can aggregate dispersed knowledge about scarcity and expectations, though doing so requires institutions that make relevant information observable and tradable (Hayek 1945; Grossman and Stiglitz 1980). Compute swaps would need benchmarks that participants regard as sufficiently objective for settlement.

Constructing such benchmarks presents unique challenges because computational capacity is inherently heterogeneous. A unit of computational service cannot be defined

solely by processor speed or nominal hardware characteristics. Economic value depends simultaneously upon processor architecture, memory capacity, networking bandwidth, storage performance, software compatibility, workload characteristics, latency requirements, and geographic location. Consequently, no single benchmark is likely to capture every commercially relevant dimension of computational performance.

Rather than viewing this diversity as an obstacle, it should instead be understood as characteristic of mature derivative markets serving heterogeneous underlying assets. Crude oil markets utilize multiple regional benchmarks such as Brent and West Texas Intermediate. Electricity markets employ hundreds of location-specific pricing nodes reflecting transmission constraints and regional demand conditions. Freight derivatives reference numerous shipping routes, vessel classes, and contract maturities. Computational finance is likely to evolve similarly, with multiple complementary benchmarks serving different segments of the computational ecosystem.

Several categories of benchmark appear especially promising. Market-based indices tracking published cloud service pricing could support contracts designed to hedge operational expenditures. Hardware performance indices could facilitate technology-transition swaps by measuring improvements across successive processor generations. Composite benchmarks incorporating energy efficiency, computational throughput, and utilization could support contracts linked to the economics of infrastructure ownership. Over time, competition among benchmark providers would likely encourage increasing transparency, methodological refinement, and market acceptance.

6.2 Settlement Mechanisms

Given the heterogeneous nature of computational services, early compute swaps would almost certainly settle financially rather than through physical delivery. Counterparties would continue procuring computational resources through their existing commercial relationships while derivative contracts independently exchanged cash flows reflecting differences between contractual reference values and observed benchmark prices or performance measures.

Financial settlement offers several advantages. It avoids the logistical complexities associated with transferring computational resources between counterparties while allowing organizations to hedge financial exposures without disrupting operational procurement decisions. Similar arrangements already dominate interest rate swaps, equity index derivatives, freight agreements, and numerous electricity contracts, demonstrating that economically meaningful risk transfer does not require physical delivery of the underlying asset.

Settlement frequency would depend upon the underlying exposure. Contracts linked to cloud service pricing might settle monthly or quarterly, while technology-transition swaps could reference annual benchmark revisions. More specialized agreements, such as utilization or operating margin swaps, might settle according to financial reporting periods or other commercially relevant intervals.

6.3 Documentation, Collateral, and Credit Risk

Like most over-the-counter derivatives, compute swaps would require standardized legal documentation to facilitate market participation while preserving contractual flexibility. Existing International Swaps and Derivatives Association (ISDA) Master Agreements provide a natural institutional foundation upon which computational derivatives could develop. New product definitions would specify benchmark methodologies, observation periods, settlement procedures, hardware classifications, and other computationally specific contractual terms while relying upon well-established legal frameworks governing payment obligations, default events, and dispute resolution.

Counterparty credit risk would initially remain one of the principal considerations in market development. Bilateral collateral agreements, variation margin, independent amounts, and credit support annexes would likely resemble those already employed throughout over-the-counter interest rate, commodity, and foreign exchange markets. Because many early participants would be sophisticated institutional counterparties, existing collateral management practices could largely be adapted rather than reinvented.

Should computational derivatives eventually achieve substantial trading volume and contractual standardization, centralized clearing could emerge as a later stage of institutional evolution. Clearinghouses reduce systemic credit exposure by interposing themselves between counterparties, standardizing margin requirements, and guaranteeing contractual performance. Financial history suggests, however, that centralized clearing generally follows market maturity rather than initiating it. Accordingly, bilateral markets are likely to dominate the early development of computational finance.

6.4 Market Participants

Several categories of market participants appear likely to contribute to the development of compute swaps, each motivated by distinct economic objectives.

Commercial consumers - including artificial intelligence developers, pharmaceutical companies, engineering firms, financial institutions, manufacturers, universities, and research laboratories - would primarily seek to stabilize future computational expenditures and reduce uncertainty surrounding long-term projects.

Infrastructure providers - including hyperscale cloud operators, data center owners, specialized GPU providers, and computational service vendors - would seek to stabilize revenues, utilization rates, operating margins, and investment returns associated with increasingly capital-intensive infrastructure.

Financial intermediaries would provide liquidity, facilitate price discovery, and match counterparties possessing complementary exposures. As computational markets mature, banks, commodity merchants, and specialized trading firms could assume roles analogous to those they currently occupy in electricity, freight, and commodity derivatives.

Institutional investors may eventually participate as computational risk becomes recognized as an independent asset class exhibiting return characteristics distinct from traditional equities, bonds, commodities, and real estate. Such participation would further enhance market liquidity while improving the allocation of computational risk throughout the broader financial system.

6.5 Institutional Evolution

Derivative markets rarely settle immediately into a final institutional form. Contract terms, benchmark methodologies, collateral practices, and trading conventions change as participants learn which exposures are common enough to standardize and which remain idiosyncratic. Compute should be expected to evolve in the same way, but not along a single path from bilateral contracts to exchanges.

The announcements of exchange-traded GPU futures already show one branch of that evolution. At the same time, better benchmarks may make some bilateral exposures easier to document and value rather than eliminating them. Standardization can therefore deepen both markets: it can move common price risk onto exchanges while giving customized contracts more credible reference points for basis, spread, capacity, or technology-related settlements.

The institutional question is thus empirical. If exchange benchmarks explain most of the variation in users' actual costs, demand for bespoke hedges may remain small. If large and persistent residual exposures survive, dealers and commercial counterparties have an incentive to design contracts around them. The relative importance of futures and swaps should reveal how much economically meaningful heterogeneity remains after benchmark construction.

7. Economic Implications

If compute swaps develop, their first order economic function would be risk allocation rather than the creation of a new asset class. Firms making long-lived investments can

separate some unwanted computational exposure from the operating decisions that generate it. That is the familiar corporate rationale for hedging: not the elimination of risk as such, but the protection of investment and financing plans from risks the firm has no comparative advantage in bearing (Froot, Scharfstein, and Stein 1993; Geczy, Minton, and Schrand 1997).

Price discovery is a possible second-order benefit, but it should not be assumed. Forward prices can reveal expectations about future scarcity, supply, demand, electricity costs, semiconductor availability, and infrastructure investment. Whether a compute derivative adds information beyond spot markets and long-term service contracts will depend on liquidity and participation. Thin bilateral markets may transfer risk effectively while contributing little public information; liquid standardized contracts may do more for observable price discovery.

Capital formation is similarly conditional. Data centers and specialized computational infrastructure require large, irreversible commitments to hardware, networking, cooling, and power. A hedge that reduces a material cash flow exposure could lower financing frictions at the margin. But the effect depends on hedge quality, collateral requirements, counterparty credit, and the possibility that derivative losses create liquidity demands of their own. The existence of a contract is not equivalent to a lower cost of capital.

The welfare implications are therefore ambiguous enough to deserve empirical treatment. Derivatives can improve risk sharing and information aggregation, but they can also introduce trading costs, model risk, counterparty exposure, and incentives to trade on imperfect benchmarks. Newbery (1987), for example, shows that the relationship between futures trading and spot market stability is not mechanically benign. Compute derivatives should be evaluated by the same standard rather than presumed beneficial because they extend financial markets into a new domain.

Several boundary cases matter. A large cloud customer may obtain more effective protection from a multi-year service agreement than from a financial swap. A vertically integrated AI firm may choose to own capacity. A small user may find the fixed costs of an OTC hedge prohibitive. Some technology-transition or availability exposures may be too difficult to define objectively for reliable cash settlement. These cases do not defeat the institutional proposition; they define the conditions under which it can be tested.

The most useful empirical questions are consequently concrete: how closely do exchange benchmarks track users realized compute costs; how large and persistent are regional and hardware-specific bases; which firms have exposures large enough to justify collateralized bilateral contracts; and do benchmark improvements shift activity from bespoke

agreements toward standardized markets? Those questions connect the proposed contracts to observable market behavior rather than treating financialization as an inevitable endpoint.

The framework yields a simple set of observable implications. Exchange trading should be most successful where prices are transparent, performance is comparable, participation is broad, and basis risk is manageable. Customized contracts should be more valuable where exposures are measurable but multidimensional or firm specific. Where measurement is weak, contracting costs are high, or commercial arrangements already provide an adequate hedge, neither futures nor swaps need dominate. Changes in benchmark quality should move these boundaries over time.

8. Conclusion

Compute is becoming expensive enough, important enough, and sufficiently forward-committed to generate demand for risk transfer. The interesting question is no longer whether every computational exposure can be represented by one financial contract. It is which parts of the exposure can be standardized without leaving so much basis risk that the hedge loses its purpose.

That distinction separates this paper from Earle (2026). The earlier analysis described a path toward benchmark-adjusted, cash settled compute futures. The present paper begins where that argument stops. CME-Silicon Data and Nodal-Compute Desk contracts show that rental price risk for specified GPU classes can be standardized. They do not make location, availability, utilization, technology transitions, workload requirements, or operating margins disappear. Those residuals are the potential domain of customized derivatives.

Swaps are therefore neither an inevitable precursor to futures nor a more advanced alternative to them. They are one contractual response to a particular combination of conditions: a material exposure, a settlement variable that can be observed, and enough mismatch with standardized benchmarks to justify the extra costs of bilateral contracting. If those conditions are absent, firms may use futures, commercial contracts, vertical integration, diversification, or no hedge at all.

This conditional formulation makes the argument testable. As compute markets develop, common exposures should migrate toward standardized contracts when benchmark quality improves and basis risk falls. Bespoke contracts should persist where economically important differences remain measurable but cannot be compressed into a common index. Risks that cannot be measured reliably may never support a derivative market. The

changing boundary among those three outcomes is more informative than any simple prediction that compute will move from one contractual form to another.

Computational capacity is useful in this respect because technological change continuously alters both the resource being priced and the risks surrounding it. New hardware can reduce one basis while creating another; new benchmarks can make a formerly bespoke exposure tradable; new commercial arrangements can make a financial hedge unnecessary. The market structure that emerges will therefore be endogenous to the economics of compute itself. That is the central institutional claim, and it is one that the next several years of market development should allow researchers to observe directly.

References

- Arrow, Kenneth J. 1964. "The Role of Securities in the Optimal Allocation of Risk-Bearing." *Review of Economic Studies* 31(2): 91-96.
- Bapna, Ravi, Sanjukta Das, Robert Day, Robert Garfinkel, and Jan Stallaert. 2011. "A Clock-and-Offer Auction Market for Grid Resources When Bidders Face Stochastic Computational Needs." *INFORMS Journal on Computing* 23(4): 630-647.
- Bapna, Ravi, Sanjukta Das, Robert Garfinkel, and Jan Stallaert. 2008. "A Market Design for Grid Computing." *INFORMS Journal on Computing* 20(1): 100-111.
- Black, Fischer. 1976. "The Pricing of Commodity Contracts." *Journal of Financial Economics* 3(1-2): 167-179.
- Brennan, Michael J. 1958. "The Supply of Storage." *American Economic Review* 48(1): 50-72.
- CME Group. 2026a. "CME Group and Silicon Data to Launch Compute Futures on October 5 to Unlock New Way to Hedge AI Risks." August 11, 2026.
- CME Group. 2026b. "Compute Futures." Contract specifications for Silicon Data H100 Rental Index Futures and Silicon Data B200 Rental Index Futures.
- Coase, Ronald H. 1937. "The Nature of the Firm." *Economica* 4(16): 386-405.
- Du, Anna Ye, Sanjukta Das, and R. Ramesh. 2013. "Efficient Risk Hedging by Dynamic Forward Pricing: A Study in Cloud Computing." *INFORMS Journal on Computing* 25(4): 625-642. <https://doi.org/10.1287/ijoc.1120.0526>.
- Earle, Peter C. 2026. "The Financialization of Compute: Forward Markets, Futures, and the Economics of Computational Capacity." SSRN Working Paper, July 2, 2026. <https://doi.org/10.2139/ssrn.7039518>.

- Fama, Eugene F., and Kenneth R. French. 1987. "Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage." *Journal of Business* 60(1): 55-73.
- Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein. 1993. "Risk Management: Coordinating Corporate Investment and Financing Policies." *Journal of Finance* 48(5): 1629-1658.
- Geczy, Christopher, Bernadette A. Minton, and Catherine Schrand. 1997. "Why Firms Use Currency Derivatives." *Journal of Finance* 52(4): 1323-1354.
- Grossman, Sanford J., and Joseph E. Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70(3): 393-408.
- Hayek, F. A. 1945. "The Use of Knowledge in Society." *American Economic Review* 35(4): 519-530.
- Hirshleifer, David. 1988. "Residual Risk, Trading Costs, and Commodity Futures Risk Premia." *Review of Financial Studies* 1(2): 173-193.
- Hull, John C. 2022. *Options, Futures, and Other Derivatives*. 11th ed. Pearson.
- Jensen, Michael C., and William H. Meckling. 1976. "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure." *Journal of Financial Economics* 3(4): 305-360.
- Keynes, John Maynard. 1930. *A Treatise on Money, Vol. II: The Applied Theory of Money*. London: Macmillan.
- Newbery, David M. 1987. "When Do Futures Destabilize Spot Prices?" *International Economic Review* 28(2): 291-297.
- Nodal Exchange. 2026. "Nodal Exchange and Compute Desk to Launch Compute Futures Contracts." September 3, 2026.
- Stulz, Rene M. 1984. "Optimal Hedging Policies." *Journal of Financial and Quantitative Analysis* 19(2): 127-140.
- Telser, Lester G. 1981. "Why There Are Organized Futures Markets." *Journal of Law and Economics* 24(1): 1-22.
- Williamson, Oliver E. 1979. "Transaction-Cost Economics: The Governance of Contractual Relations." *Journal of Law and Economics* 22(2): 233-261.
- Working, Holbrook. 1949. "The Theory of Price of Storage." *American Economic Review* 39(6): 1254-1262.