JUNE 2026 ANALYSIS

# Energy Cost Sensitivity & Bank Stress Testing

**THE REGULATORY SIGNAL IS CLEAR, THE DATA GAP IS REAL, AND THE WINDOW FOR FIRST-MOVER RISK AVOIDANCE IS NOW**


## AUTHORS

**John Harrington**
Senior Director, Banking and Finance Strategy | LandGate Corp

**Ain Siddiqui**
VP, Infrastructure Products | LandGate Corp


## BUSINESS SOLUTIONS


833-782-5837
ENERGY@LANDGATE.COM

201 MILWAUKEE ST. SUITE 200
DENVER, CO 80206


## ABOUT LANDGATE

LandGate is the leading provider of data solutions for site selection, origination, development, financing, and market analysis of renewable energy and infrastructure projects: data centers, energy storage, solar, EV's, wind, and natural gas.

# EXECUTIVE SUMMARY

The Federal Reserve has already introduced energy commodity shocks as explicit variables in the stress testing framework for the largest U.S. banks. The trajectory of scenario design (expanding variable scope, adding new guides, and building institutional capacity through the 2023 pilot Climate Scenario Analysis) points in one direction. Energy cost sensitivity for energy-intensive asset classes is not a future regulatory concept. It is a present supervisory concern in the process of becoming standard.

This paper makes three arguments:

- Energy price shocks are already embedded in the Federal Reserve's 2026 Global Market Shock (GMS) framework, which applies directly to the largest U.S. bank holding companies, including Bank of America, Citigroup, JPMorgan Chase, Goldman Sachs, Morgan Stanley, Barclays US, and Wells Fargo.
- The Fed's trajectory on scenario design (expanding variable guides, conducting exploratory climate scenario analysis, and moving toward more granular sectoral sensitivity) establishes a clear regulatory direction even in the absence of a formal energy-sector mandate today.
- The data infrastructure required to respond credibly to this regulatory direction (independent, asset-level energy cost intelligence covering grid capacity, nodal power pricing, interconnection feasibility, and project-level data) does not exist in current bank underwriting or credit risk models. LandGate is a steward of this data and observes ample reasoning for these regulatory trends.

The window for first-mover advantage and emerging risk mitigation is real, current, and narrow. The banks that build energy cost modeling infrastructure before the next rulemaking cycle will set the examination standard and likely avoid, or at least minimize, losses from these emerging risks. Those that wait will be required to reconstruct it under examiner pressure.

## The Global Market Shock Is Not Hypothetical: It Applies Now

The Federal Reserve's annual supervisory stress test applies to U.S. bank holding companies with $100 billion or more in total consolidated assets. For large banks with significant trading activity, the test includes a Global Market Shock (GMS) component, a calibrated set of hypothetical shocks to a broad range of financial risk factors, designed to simulate acute market distress and heightened uncertainty.

The 2026 GMS documentation explicitly identifies energy commodities as scenario variables. The energy index incorporated into the GMS covers fuel-based commodities including electricity, crude oil, coal, natural gas, and propane.[1] This is not a peripheral or ancillary variable. Energy commodity prices move in direct correlation with the credit quality of energy and energy-intensive borrowers, with project finance cash flow assumptions, and with the operating cost structures of data center and infrastructure assets. The shock is modeled precisely because those correlations are consequential at scale.

Bank of America, JPMorgan Chase, Citigroup, Goldman Sachs, Morgan Stanley, Wells Fargo, and Barclays U.S. are explicitly named as GMS-subject institutions in the Federal Reserve's 2026 stress test documentation. For these institutions, energy commodity shocks are an active stress testing reality, not a pending proposal. The question is whether the credit models behind the GMS exposure reflect the same sensitivity. At most institutions, they do not.

The practical consequence is a basic risk problem. Any large bank with material exposure to energy-adjacent lending (project finance, data center infrastructure, multifamily or heavy industrial commercial real estate, or C&I lending to energy-intensive industries) carries a structural mismatch between the energy commodity shocks modeled in the GMS and the energy cost assumptions embedded in its underlying credit risk models. That gap is, at most institutions, currently unaddressed and unmeasured.

## The 2026 Framework Expansion: More Variables, More Granularity

The 2026 stress test cycle introduced meaningful structural changes to how the Federal Reserve designs and communicates its scenarios. For the first time, the Fed published proposed scenarios for public comment before finalizing them, a deliberate signal that scenario design is becoming more durable, rule-based, and open to external scrutiny. The cycle also produced a new Scenario Design Policy Statement and a new small-scale macroeconomic model governing the generation of scenario variable paths.[2]

The most significant development for this analysis is the expansion of "guides," the Fed's documented ranges within which it calibrates the paths of individual scenario variables. The 2026 cycle added new guides for commercial real estate prices, equity prices, VIX, five- and ten-year Treasury yields, BBB corporate spreads, and mortgage spreads.[3] Each addition reflects the same logic: a risk category that is material enough, concentrated enough in bank portfolios, and difficult enough to model with prior-generation tools to warrant its own formal analytical framework.

That logic applies with equal or greater force to energy cost sensitivity. The GMS already stresses energy commodity prices at the trading book level. The macroeconomic scenario now formally guides CRE, credit spreads, and equity valuations. The remaining gap, energy cost sensitivity at the credit book level across energy and energy-intensive asset classes, is precisely the space the next expansion cycle addresses. The variable universe is not static. It expands toward material risk.

## The 2023 Climate Scenario Analysis Exercise

In September 2022, the Federal Reserve announced that six of the nation's largest banks (Bank of America, Citigroup, Goldman Sachs, JPMorgan Chase, Morgan Stanley, and Wells Fargo) would participate in a pilot Climate Scenario Analysis (CSA) exercise.[4] The exercise launched in early 2023 and concluded at year-end. Its stated purpose was capacity-building: the Fed was explicit that the pilot carried no capital or supervisory consequences.

The methodology it developed is the important part. The CSA analyzed how energy transition pathways and physical climate scenarios affect specific loan portfolios, geographic exposures, and asset-level cash flows. That is precisely the analytical framework that energy cost sensitivity analysis at the credit book level requires. The Fed was not investigating whether such analysis was conceptually possible. It was building the institutional knowledge and supervisory infrastructure to make it, eventually, standard.

The six institutions named in the 2023 CSA pilot are all also subject to the Global Market Shock in the 2026 stress test. The overlap is deliberate. The Federal Reserve has been building energy and climate risk analytical capacity in parallel, across both its supervisory and stress testing programs, using the same set of systemically important institutions as the development ground.

The current regulatory environment requires a note of precision. U.S. banking regulators have withdrawn from the Network for Greening the Financial System, rescinded interagency climate risk guidance, and shifted away from climate-specific supervisory mandates.5 The argument here is not premised on a climate mandate arriving. It is premised on a more tangible risk: energy cost sensitivity directly affects credit risk metrics, allowing lenders to improve underwriting accuracy and credit risk modeling. Those benefits are independent of any particular regulatory posture, and they are directly relevant to the stress testing obligations that already exist today.

## The Historical Pattern: Stress Variables Expand Toward Material Risk

The history of the Federal Reserve's stress testing framework is a consistent record of expanding variable scope in response to emerging, portfolio-concentrated risks. Commercial real estate received heightened scenario focus after the NYCB episode and the broader stress in office and rent-regulated property markets through 2023 and 2024. Corporate bond spreads received more granular treatment after the COVID credit dislocation exposed the inadequacy of index-level assumptions. In each case, the pattern was the same: the Fed identified a class of risk that was material, concentrated in bank portfolios, and inadequately captured by existing scenario tools, and then built formal analytical infrastructure around it.

Energy and infrastructure exposure now fits that description exactly. Consider the convergence of three factors:

- U.S. bank project finance and clean energy lending has scaled rapidly over the past decade, with the major bank holding companies having collectively deployed hundreds of billions of dollars across solar, wind, storage, and grid infrastructure transactions.
- Data center lending has become one of the fastest-growing segments of infrastructure finance, with energy availability and grid capacity functioning as primary determinants of asset performance and debt coverage. From our position as curators of U.S. infrastructure data, we can see the demand side of this equation (data centers) massively outpacing the supply side (energy projects).
- The Federal Reserve's own GMS framework already stresses energy commodity prices at the trading book level, creating a recognized, documented supervisory sensitivity that has not yet been connected to the credit models behind the underlying loan portfolios.

The gap between where the GMS currently operates and where credit-side scenario analysis needs to go is the space that energy cost sensitivity analysis fills. The direction of travel is not in question. The only open question is which institutions will be early adopters and minimize or altogether avoid the losses these risks pose to their portfolios, and which ones will have to do so retroactively after a major event, or events, of default.

## The Project Finance Model and the Grid Capacity Blind Spot

Consider how a project finance bank underwrites a utility-scale solar facility today. The model captures construction cost, technology performance curves, PPA contract terms, sponsor credit quality, and debt service coverage across a range of production scenarios. What it does not capture, and what no standard underwriting model currently captures, is the grid infrastructure context: whether transmission capacity to deliver the project's output actually exists, where the project sits in the interconnection queue, how applicable power prices, hub prices, and location-specific basis, may behave over the asset's operational life, and what competing load demands on the local grid look like over a short, mid, and long term horizon.

Grid capacity is the binding constraint on renewable energy and energy-intensive development projects across large portions of the United States today. Projects that appear financially sound at underwriting can be rendered economically marginal by years-long interconnection delays, transmission constraints, deliverability limits, or location-specific basis differentials between project-level prices and broader market benchmarks that erode profitability assumptions. None of these risks appear in standard underwriting models, not because they are obscure, but because the data infrastructure to model them at the asset level has not been viewed as relevant until recently.

---

The magnitude of the interconnection problem is visible in LandGate's own grid data. Across a sample of PJM substations in Virginia, 106 interconnection requests representing 7,713 MW of nameplate capacity had been filed. Of those, 65% had been withdrawn before reaching commercial operation, and only 19% remained active. In other words, roughly two of every three projects that a credit model might treat as a comparable, or as eventual grid supply, never reach the grid at all. A capacity factor or supply assumption that ignores queue attrition of this magnitude is not conservative. It is structurally optimistic.

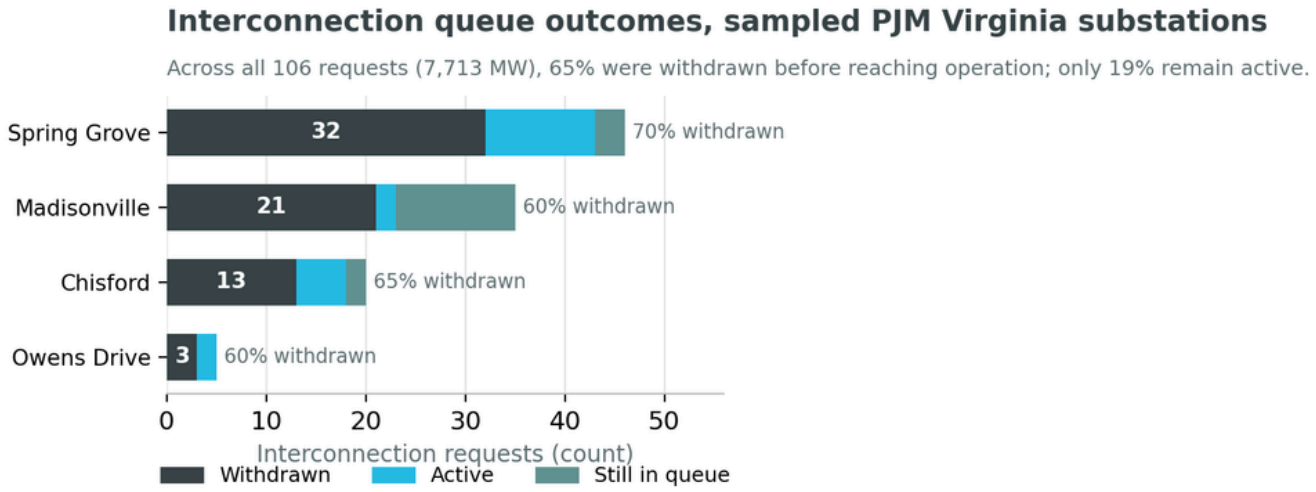


LandGate's proprietary data covers U.S. grid capacity mapping, historical and forecast LMPs, settlement point prices, and hub/zone price benchmarks, interconnection queue positions, project-level renewable energy development data, and data center infrastructure intelligence, providing the independent, asset-level energy intelligence layer that current bank models lack.

## The Data Center Sector: Credit Exposure on Both Sides of the Balance Sheet

The data center sector presents a compounded version of this problem. The largest U.S. banks carry data center exposure simultaneously in two directions: as lenders to third-party data center development and infrastructure projects, and as direct operators of significant owned technology infrastructure. For institutions with both exposures, the grid intelligence gap affects credit underwriting and internal capital planning at the same time.

Bank of America, as an instructive example, deploys billions annually in technology investment, anchored in a private cloud-first infrastructure strategy. The decisions embedded in that investment (site selection, long-term energy cost projections, grid capacity planning at owned facilities) are being made using internal projections and utility-provided data. Neither provides the independent, third-party grid intelligence needed to model operating cost trajectories accurately at the asset level, in the context of locality and competing projects that will impact the landscape they operate in.

On the lending side, the problem is equally acute. Data center finance has grown into one of the most capital-intensive segments of infrastructure lending, and energy cost is not a secondary operating item in these transactions. It is frequently the dominant driver, either directly or indirectly, of operating cash flow, representing 40% or more of total operating expenses in hyperscale and enterprise-scale facilities. Even a 20% increase in power prices, or a widening in location-specific basis, can materially impair debt service coverage for a project that appeared well-structured at origination.

LandGate observes energy cost volatility directly at the nodal level, and the dispersion is severe even outside the most congested corridors. Across a sample of PJM pricing nodes in Virginia in June 2026, hourly day-ahead LMPs ranged from -$0.84/MWh to $194/MWh, while the average across sampled nodes and intervals was roughly $41/MWh. Single-node peak prices ran 3.0 to 4.2 times the node's own monthly average, and one of the five sampled nodes priced negative at least once within the month. In the most volatile markets, primarily the northeastern U.S., the monthly swings are wider still. This is the asset-level reality that a smoothed regional price assumption cannot represent.

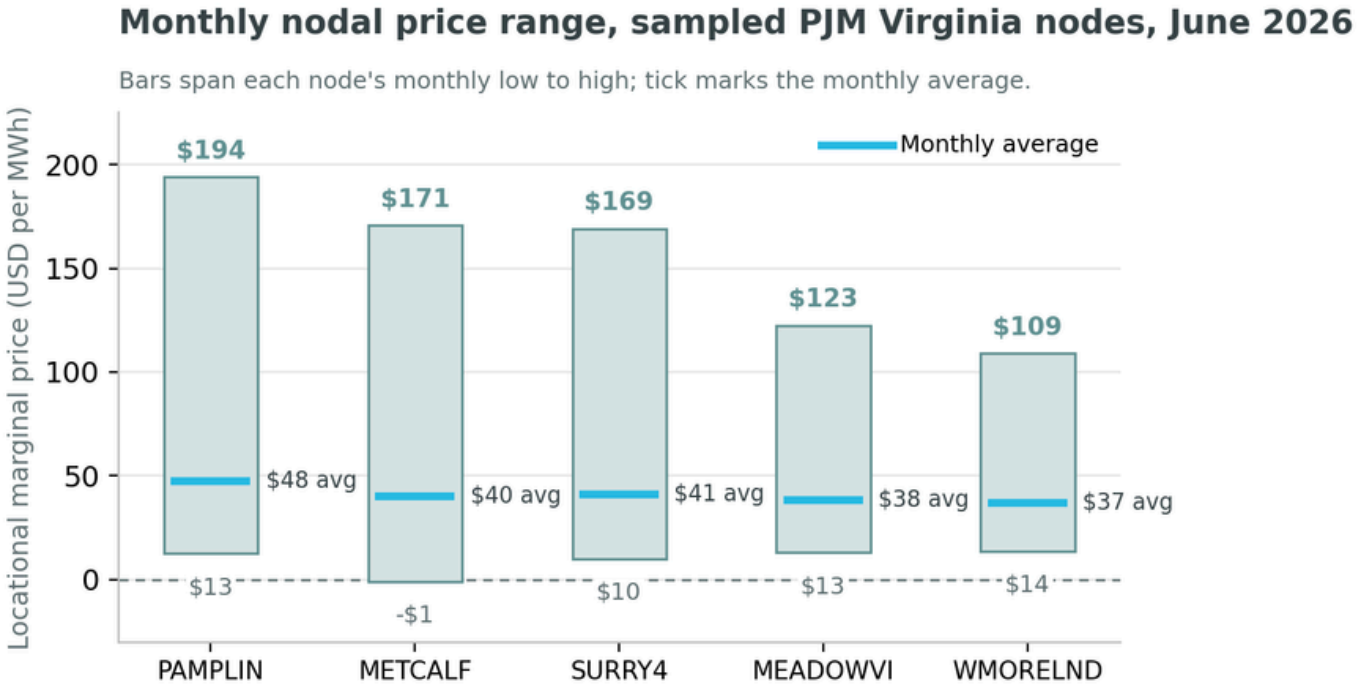

Some lenders say they are covered because they only underwrite NNN (Triple Net) leases on data centers. Hyperscale tenants represent some of the most sophisticated and profitable tech companies in the history of the world. Lenders who think they are immune should review the exit clauses within those leases. The penalties are steep, but not sufficient to repay exposure. This is not a tail-risk scenario. It is a realistic outcome in congested grid markets where nodal basis differentials are becoming increasingly severe. The gap in managing this risk is not a modeling sophistication problem. It is a data awareness and availability problem.

**Asset-level signals current bank models miss**

*LandGate grid intelligence, sampled PJM substations and nodes in Virginia, data pulled June 2026.*

| Signal | What LandGate observes in the sample |
|---|---|
| **Interconnection queue attrition** | 106 interconnection requests (7,713 MW nameplate) filed across the sampled substations. 65% were withdrawn before reaching operation; only 19% remained active. Roughly two of every three queued projects never reach the grid. |
| **Nodal price level** | Cross-node monthly average locational marginal price near $41/MWh, but individual nodes ranged from negative ($0.84 below zero per MWh) to $194/MWh within the single month. |
| **Intra-month price dispersion** | Single-node peak prices ran 3.0 to 4.2 times the node's own monthly average. One of five sampled nodes priced negative at least once in the month, illustrating volatility that regional averages cannot capture. |

*Sample drawn from in-service PJM substations in Virginia. Figures are illustrative of the asset-level dispersion that smoothed, regional inputs obscure, and are not a portfolio-wide forecast.*

## Credit Risk Modeling and the GMS-to-Credit-Book Disconnect

The Federal Reserve's supervisory stress tests require large banks to model loan losses, revenues, and capital adequacy across hypothetical macroeconomic scenarios. For banks with significant energy and infrastructure lending, the accuracy of those models depends on whether energy cost variables (nodal power pricing, grid availability, interconnection outcomes) are reflected in the underlying credit risk models at the asset level.

Some credit risk models for project finance, infrastructure lending, and commercial real estate mildly address this using EIA-published or utility-provided energy price assumptions that reflect national or regional averages. They do not capture nodal-level pricing differentials, which deviate from regional averages by 30% to 50% at specific locations with regularity, a gap the Virginia node sample above illustrates directly. They do not capture interconnection queue status, which directly determines whether a project's assumed capacity factor is operationally achievable, and which, as the queue data shows, eliminates most projects before they ever reach the grid. And they do not capture grid capacity constraints, which have become the primary performance variable for energy-intensive assets in congested markets.

The result is a structural mismatch. The GMS stresses energy commodity prices at the index level. The credit models behind those same portfolios use smoothed, averaged energy cost inputs that bear little relationship to asset-level reality. Closing that gap is not a compliance exercise. It is a fundamental improvement in the accuracy of credit risk assessment for one of the fastest-growing and most energy-sensitive categories of bank lending.

# IV. THE FIRST-MOVER CASE:
## TIMING, POSITIONING, AND COMPETITIVE ADVANTAGE

Regulatory frameworks rarely arrive without warning. They are preceded by a visible sequence: pilot exercises, public comment processes, documented examiner concern, and expanding variable guidance. Institutions that pay attention to that sequence position themselves to shape the standard. Those that wait find themselves responding to it.

The 2024 and 2025 heightening of CRE stress scenarios did not catch attentive institutions by surprise. The escalating Fed scrutiny of office concentration and rent-regulated property exposure had been documented for two years before the 2026 scenario design formally incorporated CRE price guides. The institutions that built their CRE stress testing infrastructure ahead of that formalization presented to examiners as sophisticated and proactive. The institutions that had not presented very differently.

The energy cost sensitivity trajectory is at roughly the same stage today that CRE stress was in 2022. The signal is clear. The formal requirement has not yet arrived. The analytical infrastructure required to respond credibly takes time to build.

In the supervisory relationship, the distinction between anticipating a requirement and responding to examiner pressure is not cosmetic. It shapes examination ratings, model validation outcomes, the quality of the ongoing dialogue with the Fed, and, over time, the regulatory capital treatment applied to the relevant portfolios.

The competitive dimension compounds the regulatory argument. The largest U.S. banks are competing for the same project finance mandates, infrastructure lending relationships, and data center financing transactions. In deals where energy infrastructure context is central to underwriting quality, the institution that can demonstrate independent grid intelligence, showing sponsors and co-lenders that it understands grid capacity constraints, nodal pricing dynamics, and interconnection risk at the specific project level, is in a structurally stronger position to win the mandate and retain the relationship.

In project finance and infrastructure lending, analytical differentiation at origination translates directly into deal economics: better structure, tighter pricing, stronger syndication execution, and deeper sponsor relationships. The institutions that build energy cost modeling capability now are building a compounding advantage in one of the most active growth segments of U.S. bank lending. The institutions that wait are ceding that ground incrementally, one transaction at a time.

# V. CONCLUSION

The regulatory direction is unmistakable, the analytical gap is real and present, and the first-mover advantage for institutions that act ahead of that direction is material.

Three things are simultaneously and independently true:

- Energy commodity shocks are already embedded in the Federal Reserve's stress testing framework for the largest U.S. banks: a present, documented supervisory sensitivity to energy cost variables that has not yet been connected to underlying credit risk models.
- The Fed's trajectory on scenario design is directionally clear: more variables, more granularity, more formal guidance, and a documented history of building supervisory infrastructure around material risk concentrations before formalizing them as requirements.
- The data infrastructure required to respond to this direction, independent asset-level energy cost intelligence at the project and portfolio level, does not exist in current bank models, and the window to build it proactively is finite.

LandGate's energy and infrastructure intelligence platform provides the data layer that bridges the gap between the energy commodity shocks the U.S. Federal Reserve already models and the credit risk framework that needs to reflect them. The case for acting now is not regulatory compliance for its own sake. It is better underwriting, more accurate credit risk assessment, and more credible stress testing, with the regulatory trajectory providing an additional rationale that only strengthens over time.

## Sources


[1] Board of Governors of the Federal Reserve System. Supervisory Stress Test Documentation: Final 2026 Global Market Shock Component. February 2026. federalreserve.gov/supervisionreg/files/2026-final-gms-model.pdf

[2] Board of Governors of the Federal Reserve System. 2026 Stress Test Scenarios. February 2026. federalreserve.gov/publications/files/2026-final-supervisory-stress-test-scenarios-20260204.pdf

[3] Bank Policy Institute. "Federal Reserve Proposes 2026 Stress Test Scenarios." November 19, 2025. bpi.com/federal-reserve-proposes-2026-stress-test-scenarios

[4] Board of Governors of the Federal Reserve System. Press Release: "Federal Reserve Board announces that six of the nation's largest banks will participate in a pilot climate scenario analysis exercise." September 29, 2022. federalreserve.gov/newsevents/pressreleases/other20220929a.htm

[5] Resources for the Future. "The Evolving View of Climate-Related Financial Risks in the US Financial Sector." December 15, 2025. resources.org/common-resources/the-evolving-view-of-climate-related-financial-risks-in-the-us-financial-sector

## About LandGate

LandGate is the leading energy and infrastructure intelligence platform for the U.S. grid, power market, and data center ecosystem. LandGate's proprietary data covers U.S. grid capacity mapping, nodal power pricing (historical and forecast), interconnection queue analysis, renewable energy project-level data, and data center infrastructure intelligence. The platform gives banks, project finance institutions, infrastructure investors, and energy developers the independent, asset-level data layer required to underwrite energy-intensive assets with precision, and to model energy cost sensitivity at the scale and granularity that current regulatory trajectories demand.

## About the Authors

J**ohn Harrington** leads Banking and Finance Strategy at LandGate Corp, where he oversees energy and infrastructure intelligence use cases for major financial institutions. His work is focused on translating LandGate's proprietary grid, power, and data center data into institutional use cases for bank underwriting, credit risk modeling, stress testing, and investing.

Prior to LandGate, John held advisory and credit roles across the U.S. banking sector. At Moody's, he supported regional bank clients in a variety of ways, giving him direct exposure to how large institutions approach credit risk modeling, regulatory stress testing, and model validation. During this time, he was on the leading edge of the shiftings within CRE stress testing regulatory guidance as well as adjusted ALM processes following the SVB liquidity crisis. Before Moody's, he held credit risk roles at Fifth Third Bank and First Bank, where his work was concentrated in commercial and industrial underwriting and complex credit analysis. During his time at Fifth Third, John spent time as the lead underwriter for the Renewable Energy team.

John holds a degree in Commercial Banking from the McColl School of Business at Queens University of Charlotte. He is still based in Charlotte, North Carolina.

**Ain Siddiqui** is VP, Infrastructure Products at LandGate, where he works on power price forecasting, data center infrastructure, algorithmic due diligence, and capacity modeling. Over his three years at LandGate, his work has centered on turning complex grid, power market, and infrastructure datasets into practical intelligence for developers, investors, and financial institutions.

Ain holds a bachelor's degree in Mathematics from UCLA and is pursuing a master's degree in Data Science at Columbia University.