Your Solar Investment Committee Is Debating the Wrong Variables

A sensitivity analysis of 15 commonly negotiated inputs reveals that cost of debt, PPA escalators, and construction timelines barely affect solar project economics. Three other variables dominate. Knowing which ones changes how organizations screen, price, and approve clean energy deals.

Kshitiz Raj
Energy Finance Specialist, Redaptive Sustainability Services, Boulder, Colorado
Master's from Duke University, Durham, North Carolina

In 2025, corporations contracted 55.9 gigawatts of clean power through purchase agreements, with Amazon, Meta, Google, and Microsoft responsible for nearly half of global activity.[1] Behind every one of those deals sat a financial model. And in virtually every investment committee meeting that approved those models, the conversation spent disproportionate time on the wrong inputs.

That is not speculation. It is a testable claim, and we tested it.

Using a standard solar project finance model with typical U.S. debt structures and tax treatment, we measured how sensitive three critical investment outcomes are to each of 15 commonly negotiated inputs.[2] The finding that should change how organizations evaluate clean energy investments: cost of debt, the variable that dominates most boardroom discussions about project economics, ranked fourteenth out of fifteen. PPA escalators, the negotiating lever that procurement teams and sellers spend weeks on, ranked ninth. Construction duration ranked dead last.

The three variables that actually drive project value, investor returns, and the tariff a buyer must sign are delivered capital expenditure per watt (CapEx), bankable energy production (kilowatt-hours per kilowatt of capacity), and the base-year PPA price. That hierarchy holds whether you are a developer valuing equity, a lender stress-testing loan coverage, or a corporate buyer negotiating price. And it holds across all three financial metrics that investment decisions actually rest on: equity net present value, levered internal rate of return, and the minimum PPA rate required to clear investor hurdles.

For any executive approving a solar PPA, or any board member reviewing a capital allocation to clean energy, the implication is immediate: the variables your team is likely optimizing hardest are not the variables that determine whether the deal creates value.

Why Interest Rates Barely Matter in Solar Finance

Most executives carry an intuition from corporate finance that interest rates are a primary driver of investment returns. In conventional capital budgeting, that is often true. In solar project finance, under the structures that dominate the U.S. market, it is not. The gap between intuition and reality is an example of what behavioral economists call anchoring bias: decision-makers anchor on a familiar reference point and adjust insufficiently when the underlying mechanics change.[3]

Research on capital allocation has documented this pattern extensively. McKinsey's cross-sector analysis of more than 1,500 U.S. companies found that those which least anchored their capital allocation decisions to prior-year patterns outperformed peers by nearly four percentage points in median returns.[4] A Deloitte survey found that more than 60% of finance executives were not confident in their organization's ability to make optimal capital allocation decisions, in part because cognitive biases, including anchoring, distorted how teams evaluated inputs.[5] Solar investment committees exhibit the same pattern. Teams anchor on interest rates because rates dominate conventional finance discussions, even though the underlying project structure works differently.

Here is how the structure actually works. In a typical solar PPA, debt is not sized as a fixed percentage of project cost. Instead, the lender sizes the loan so that projected cash flows cover debt payments by a minimum ratio (typically 1.35 times or more) in every year of the repayment window. This is called a debt service coverage ratio, or DSCR, constraint. When interest rates rise, the model does not simply pass the higher cost through to equity returns. Instead, it rebalances the repayment schedule to maintain coverage. The DSCR floor acts as a built-in shock absorber, dampening how much rate changes actually affect the outcomes investors care about.

Exhibit 1: Why Interest Rate Changes Are Absorbed in Solar Project Finance

Step	Fixed Amortization	DSCR-Sculpted Debt
1	**Interest rate rises 1%** Same shock applied to both structures	**Interest rate rises 1%** Same shock applied to both structures
2	**Repayment schedule unchanged** Fixed principal + higher interest = higher total payment each period	**DSCR floor triggers rebalancing** Model maintains 1.35x coverage by extending tenor and reshaping payments
3	**Higher debt service in every period** Cash available to equity shrinks proportionally to rate increase	**Debt service stays near original levels** Payments shift to later years; early distributions preserved
4	**Equity absorbs the full hit** IRR drops materially, NPV declines, required PPA rises	**Equity impact is dampened** IRR, NPV, and required PPA barely move; cost of debt ranks 14th

Note: The left column shows how a rate increase flows through a conventional fixed-amortization structure, where equity absorbs the full impact. The right column shows the same rate increase under DSCR-sculpted debt, where the repayment schedule automatically rebalances to maintain coverage, *dampening the effect on equity returns. This mechanism explains why cost of debt ranked 14th out of 15 inputs in our sensitivity analysis.*

This is why cost of debt ranked fourteenth in our analysis. It is not that interest rates are irrelevant to the broader economy. It is that the specific financing structures used in most U.S. solar PPAs contain a mechanism that limits how much rate changes affect value, returns, and price. Change the structure to a longer repayment period, fixed amortization, or looser coverage requirements, and the sensitivity to interest rates would increase. Under the structures in wide use today, it does not.

The second surprise is PPA escalators. A 2% annual price increase sounds like it compounds into real money over a 20-year contract. But the time value of money works against back-loaded revenue, and lender coverage tests concentrate on the first ten years. The escalator's contribution to equity value and investor returns is modest compared to the base-year price. If a procurement team is negotiating hard on the escalator while accepting a weak starting rate, the economics of the deal are being shaped by the wrong lever.

The Three Variables That Actually Set Solar Economics

Across all three investment metrics, the same three inputs dominate by a wide margin. Exhibit 2 visualizes the full ranking of all 15 inputs across each metric. Understanding why the top three dominate is more useful than memorizing the ranking.

Exhibit 2: Sensitivity Rankings Across Three Financial Outcomes

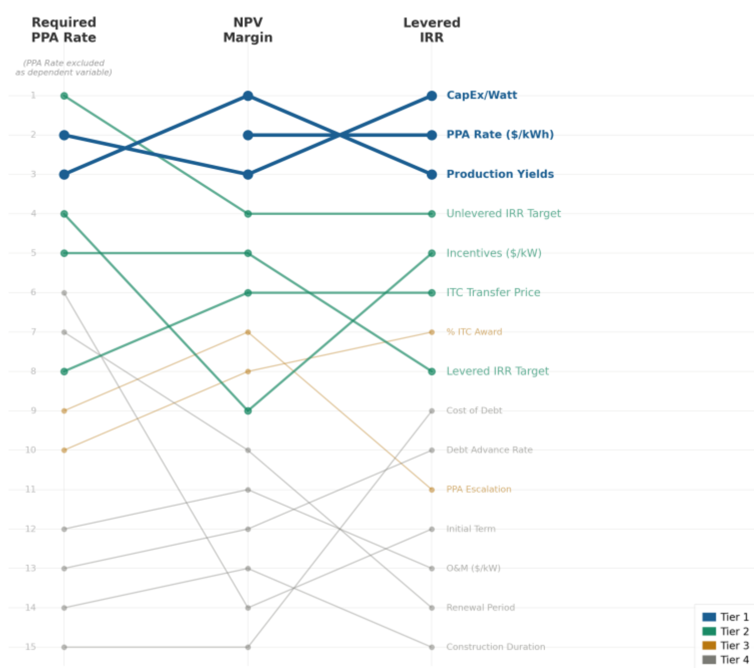

Note: Each line tracks a single input's rank across three financial outcomes. Lines that remain flat near the top indicate consistent dominance across metrics. Crossovers reveal where an input's importance shifts between outcomes. PPA Rate ($/kWh) is excluded from the Required PPA Rate column because it is the dependent variable in that analysis. A 1% symmetric shock was applied to each input; the outcome with the largest absolute percentage response ranked highest. Color groups inputs by tier based on weighted-average rank.

Source: Author's sensitivity analysis. Full model specification, base case, and equations are detailed in the methodology section.

CapEx per watt (the total cost to build each watt of capacity) sets the capital intensity of the project and shapes everything downstream: the size of the loan, interest accrued during construction, the tax basis for depreciation, and the equity check at closing. Every dollar removed from CapEx reduces required borrowing, compresses repayment, and improves early equity distributions, which disproportionately influence return calculations.

Bankable production (the kilowatt-hours per kilowatt of capacity that lenders and investors are willing to underwrite) multiplies every year's revenue over the contract horizon. It is not just a question of how much sun hits the panels. It encompasses interconnection risk (can the grid accept the energy?), curtailment (will the grid operator reduce output during oversupply?), and operating availability (will equipment perform reliably?). A project that models high output but faces a congested interconnection point will behave as a high-risk asset in practice, even if the resource looks favorable.

The base-year PPA price (the starting tariff in the contract) sets the height of the revenue line in the years when lender coverage tests bind. Unlike escalators, which shift value to later years where discounting diminishes their impact, the base rate determines whether the project meets its financial tests in the period that matters most.

When these three inputs align favorably across all three metrics, the result is something practitioners rarely get from sensitivity analysis: a single priority list that travels from the development team to the credit desk to the procurement negotiation without translation.

What This Means for Business Leaders Evaluating Clean Energy

Corporate clean energy procurement is no longer a niche sustainability function. In 2024, corporations contracted a record 68 gigawatts of power through PPAs, with solar representing roughly half.[6] The data center sector alone drove nearly 60% of U.S. corporate deals, and the number of unique buyers in the U.S. dropped by half in 2025, leaving a market increasingly dominated by hyperscalers while smaller companies pulled back amid policy uncertainty. But the procurement teams signing these contracts, and the boards approving them, are often evaluating proposals using instincts borrowed from conventional capital budgeting rather than the mechanics of project finance.

The sensitivity hierarchy reframes how buyers should evaluate competing solar proposals. When a developer offers a low price, the first question should not be whether the economics seem too good to be true. It should be whether the quoted CapEx includes full scope.

Our companion research across 164 commercial solar projects in California, Illinois, and New Jersey found that the same label of project cost can describe very different underlying content. California averaged $3.88 per watt while Illinois averaged $2.44 per watt, but much of that gap traced not to regional inefficiency but to differences in what was actually included: parking-canopy structures versus flat-roof systems, whether interconnection upgrades were bundled, and whether developer fees were captured in the reported number.[7] A procurement team evaluating bids from different developers without understanding these scope differences is comparing numbers that are not comparable. A low CapEx figure that excludes interconnection or structural work is not a better price. It is a hidden risk sitting inside the single most sensitive variable in the financial model.

The hierarchy also explains why policy shifts hit solar procurement so hard. When we modeled four executed school-district PV projects in Connecticut under paired scenarios, toggling only the Investment Tax Credit from 40% to 0% while holding everything else constant, required PPA prices rose between 42% and 60%.[8] The credit works because it reduces effective CapEx, the top-ranked variable. Remove it, and the math forces either higher tariffs or larger equity contributions at closing.

Consider what this means in practice. Denver Public Schools has deployed 8.8 megawatts of solar across 47 sites using PPAs and now generates 9.5% of its electricity from solar, saving 7% annually on energy costs.[9] Those savings exist because the ITC lowered the effective build cost, which lowered the PPA price the district pays. A 42–60% increase in required PPA would push many similar projects past the budget ceiling a school board would approve. The sensitivity hierarchy makes that risk visible in terms district administrators and board members can understand without wading through financial models: the policy lever works because it pushes on the variable that matters most.

One Priority List for Engineering, Finance, and Policy

There is a tendency in energy investment to treat cost-competitiveness and financeability as separate conversations. Engineering teams talk about levelized cost of energy. Finance teams talk about IRR and debt coverage. Policy teams talk about incentive design. Each uses different metrics and draws different conclusions about what matters most.

The sensitivity hierarchy shows these conversations converge.[10] Reduce delivered CapEx and you lower both the levelized cost and the required PPA while enhancing debt coverage and IRR. Increase net energy per kilowatt of capacity and the same thing happens. Defend a realistic base price and the project clears both the technology comparison and the investment committee.

That convergence is the most actionable takeaway for organizations trying to accelerate clean energy deployment. It means you do not need separate playbooks for the engineering team, the finance team, and the procurement team. The priority stack is the same: get the build cost right, get the energy yield right, get the contract price right. Everything else is refinement.

For corporate buyers, it clarifies that there are exactly three paths to a lower tariff without undermining bankability. Exhibit 3 maps these paths.

Exhibit 3: Three Paths to a Lower PPA Tariff

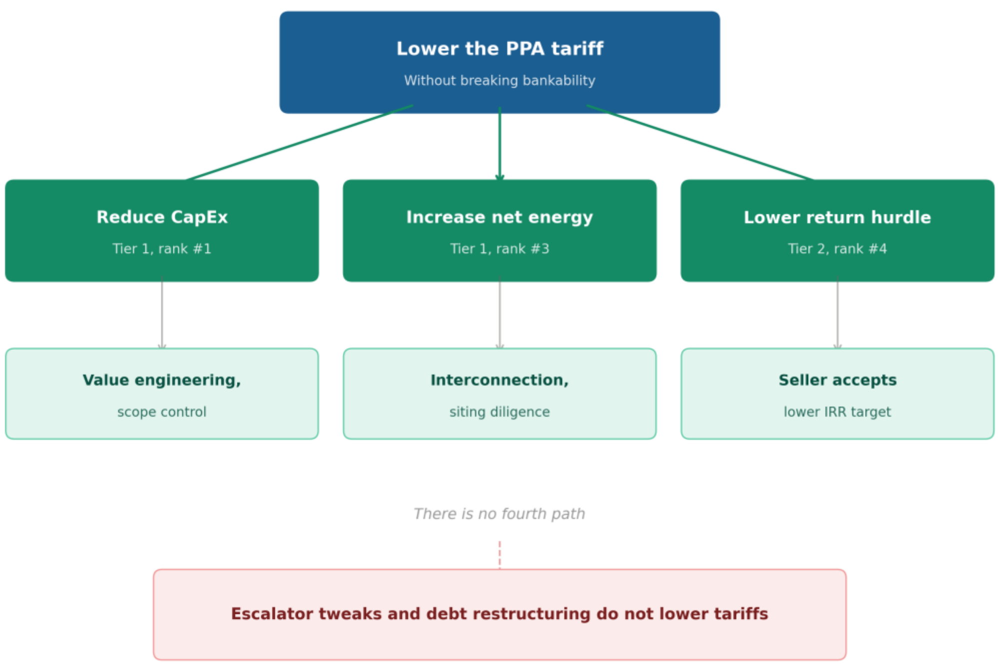

Note: The three valid paths correspond to the Tier 1 and Tier 2 variables from the sensitivity analysis. Escalator adjustments and debt restructuring (Tier 3–4 variables) do not materially reduce tariffs under standard DSCR-constrained structures.

For development teams, this justifies concentrating engineering and procurement dollars on value engineering that reduces delivered CapEx and on interconnection diligence that prevents late-stage cost surprises. For lenders, it provides a triage map: stress-test production risk and CapEx risk first, because a project that survives a 5–10% yield shortfall and a 10–15% cost overrun while keeping coverage intact is structurally sound. There is no fourth path that runs through escalator tweaking or debt restructuring.

For policymakers, the hierarchy explains why incentives matter but cannot do the whole job. Tax credits sit in Tier 2 because they close financing gaps through early-year cash, but they cannot substitute for reductions in delivered capital cost and improvements in bankable energy. The fastest megawatts unlocked come from actions that push on Tier 1 across the entire pipeline: reducing permitting friction, speeding interconnection queues, and stabilizing incentive policy so that development teams can commit capital with confidence.

When the Hierarchy Shifts

The ranking holds for the PPA-based, coverage-constrained structures that dominate U.S. solar finance today. But three conditions can shift it, and leaders should know when to rerun the analysis with different assumptions.

The first is a change in the debt product: If a project uses an interest-only mini-perm (where no principal is repaid during the initial loan period) or a bullet maturity (where the entire principal comes due at once), the shock-absorber effect of coverage-based sculpting weakens. Interest rate changes pass through to equity more directly because the repayment schedule is not being rebalanced to absorb them. In structures like these, cost of debt would climb from Tier 4 into Tier 2 or higher. Developers and lenders who routinely use mini-perms for construction-phase financing should expect a different sensitivity profile than what our base case produces.

The second is a merchant-heavy market: In independent system operators like ERCOT in Texas, where a significant share of revenue depends on wholesale spot prices rather than a fixed PPA, the production input effectively splits into two distinct risk channels: how much energy the project generates, and what that energy is worth at the moment it is delivered. In ERCOT, solar output often coincides with low or even negative wholesale prices during midday hours because of oversupply from competing solar plants. A project that looks strong on kilowatt-hours but faces chronic price compression during peak generation may find that the revenue-per-kilowatt-hour, not just the volume of kilowatt-hours, becomes the binding constraint. In that setting, production yields and price shape become co-equal drivers, and the hierarchy partially reshuffles.

The third is storage coupling: When a solar project is paired with a battery, the economics gain new revenue streams: capacity payments, energy arbitrage (storing cheap midday solar and selling it during expensive evening hours), and firming value. These streams introduce variables that do not exist in a standalone solar model, including battery degradation, augmentation schedules, and dispatch optimization. BloombergNEF tracked 5.2 gigawatts of co-located and hybrid deals in 2025, and expects this share to grow as corporate buyers seek firmer power products.[11] For these projects, the sensitivity surface changes enough that the framework should be rerun with the additional inputs included.

None of these caveats overturn the core finding for the structures in wide use. They do mean that the framework is not a fixed answer but a replicable method. The model equations, solve logic, and scenario protocol are fully documented and designed so that a practitioner can plug in their own base case and verify whether the hierarchy holds for their specific deal or market.

The United States must convert planned gigawatts into operating capacity. Every project must clear its own investment committee, its own credit desk, and its own procurement approval. The hierarchy documented here gives those decision-makers a short, defensible answer to the question they face every week: where should we focus?

If a project is strong on CapEx, production, and base price, incentives and monetization can close gaps, and the rest of the term sheet should be tuned but not obsessed over. If it is weak on those three, the lower-tier variables cannot rescue it. That distinction is not cosmetic. It delivers operational advantage: developers allocate resources efficiently, lenders stress the failure modes that matter, buyers negotiate tariffs that remain financeable, and policymakers target spending where it yields additional megawatts.

Clarity about what actually moves value, return, and required tariff is not just better analysis. It is more clean energy capacity placed in service.

RESEARCH METHODOLOGY

This study used a deterministic project finance model built in Microsoft Excel 365 with VBA macros for debt sizing and PPA solving. The model is annual and enforces DSCR-constrained amortizing debt with U.S. tax treatment (Investment Tax Credit, 5-year MACRS depreciation, OBBB treatment and, basis reduction per statute).

Three outcomes were evaluated: NPV Margin (equity NPV divided by CapEx), Levered IRR (the internal rate of return to equity on post-tax flows with DSCR-sized debt), and Required PPA Rate (the base PPA price that satisfies the sponsor hurdle while respecting all coverage tests). Local sensitivities were computed as one-at-a-time centered elasticities: a symmetric 1% shock to each input and the corresponding percentage change in the outcome. Fifteen inputs were tested, spanning system parameters (CapEx/W, production yields), contract structure (PPA rate, escalation, initial term, renewal period), equity targets (unlevered and levered IRR), debt terms (cost of debt, advance rate), tax and incentive parameters (ITC percentage, transfer price, capacity incentives), O&M cost, and construction duration.

The weighted-average master ranking was computed as the arithmetic mean of each input's rank across the three outcomes, with ties broken by the largest impact score. The impact score for each input-outcome pair was defined as the absolute percentage change in the outcome per percentage change in the input, evaluated symmetrically around the base at the 1% step.

The base case was calibrated from a documented sponsor workbook: 20-year initial PPA term with 10-year renewal option, 30% ITC with $0.91/$1 transfer price, $2.50/W all-in CapEx (including interconnection and owner costs), 1,321 kWh/kW first-year production, $0.1444/kWh nominal PPA with 2.0% annual escalation, $5.75/kW-year O&M, 6% cost of debt, 80% advance rate, and DSCR target consistent with typical lender requirements. Three automated checks ran on every scenario: balance sheet and cash waterfall reconciliation, directional comparative statics, and DSCR compliance. Global sensitivity results from Monte Carlo sweeps over calibrated ranges confirmed the local ranking.

Companion datasets include: (1) 164 commercial solar project records across California, Illinois, and New Jersey, assembled from public-facing procurement documentation and deal materials, for cost benchmarking analysis; and (2) four executed Connecticut school-sector PV projects (two front-of-meter, two behind-the-meter), modeled under paired ITC counterfactuals (40% vs. 0%) with all other inputs held constant and a 15% levered IRR target. Full model equations, solve logic, and base case parameters are available from the author.

ENDNOTES

1. BloombergNEF, "1H 2026 Corporate Energy Market Outlook: Cooling Off—Down, but Not Out" (London: BloombergNEF, February 2026). Technology giants Meta, Amazon, Google, and Microsoft accounted for 49% of global corporate PPA activity in 2025.

2. K. Raj and M. Walton, "Which Assumptions Really Set Power Purchase Prices and Returns in United States Solar Projects," Advances in Consumer Research 2, no. 5 (2025): 1960–1978.

3. D. Kahneman and A. Tversky, "Judgment Under Uncertainty: Heuristics and Biases," Science 185, no. 4157 (1974): 1124–1131. Anchoring bias leads decision-makers to rely too heavily on an initial reference point and adjust insufficiently when the underlying mechanics change.

4. D. Lovallo, A.L. Brown, D.J. Teece, and D. Bardolet, "Resource Re-Allocation Capabilities in Internal Capital Markets: The Value of Overcoming Inertia," Strategic Management Journal, cited in McKinsey, "An Analytics Approach to Debiasing Asset-Management Decisions" (New York: McKinsey & Company, December 2017). The analysis compared capital allocation and total shareholder returns over a 20-year period at more than 1,500 U.S. companies.

5. Deloitte, "Capital Allocation: Recognize Bias in Your Decision-Making," CFO Insights (New York: Deloitte, July 2023). The survey found that more than 60% of finance executives were not confident in their capital allocation abilities.

6. C. Zhu and B. Brunetti, "Data Centers Lead Global Growth in Corporate PPAs," S&P Global Commodity Insights/pv magazine USA, April 14, 2025. Solar accounted for more than half of contracted capacity in North America in 2024.

7. K. Raj and M. Walton, "Regional Disparities in Solar Photovoltaic Installation Costs: A Multi-State Analysis of Pricing Mechanisms and Scale Economies," Scientific Culture 12, no. 4 (2026): 112–130, https://doi.org/10.5281/zenodo.12426110. [Note: This companion study is currently in press at time of submission.] See also V. Ramasamy, D. Feldman, R. Fu, and M. Woodhouse, "U.S. Solar Photovoltaic System and Energy Storage Cost Benchmarks, Q1 2023," NREL/TP-7A40-87303 (Golden, Colorado: National Renewable Energy Laboratory, 2023).

8. K. Raj, "Impacts of Investment Tax Credit on Public Schools in the United States: A Paired Analysis of Four Connecticut Projects," International Journal of Environmental Sciences 11, no. 24s (2025): 3320–3331, https://doi.org/10.64252/tz4ra011.

9. J. Dorn, "Solar Is Helping Schools Save Big. Your District Could Be Next," Canary Media, October 31, 2024. See also Generation180, "Brighter Future: A Study on Solar in U.S. K–12 Schools" (Charlottesville, Virginia: Generation180, 2024).

10. For a discussion of LCOE's limitations as a financeability metric, see J. Aldersey-Williams, G. Broadbent, and P.A. Strachan, "Levelised Cost of Energy: A Theoretical Justification and Critical Assessment," Energy Policy (2019); and Lazard, "Levelized Cost of Energy Analysis, Version 17.0" (New York: Lazard, 2024). See also P. Gilman, A. Dobos, S. Janzou, et al., "SAM (System Advisor Model) General Description," NREL/TP-6A20-70414 (Golden, Colorado: National Renewable Energy Laboratory, 2018).

11. BloombergNEF, "1H 2026 Corporate Energy Market Outlook." Seven of the top 10 sellers engaged in baseload-like power contracts, including co-located solar and storage and hybrid solar and wind products.

ABOUT THE AUTHOR

Kshitiz Raj is an Investment Specialist in new clean energy technology, based in Boulder, CO. He has structured and underwritten solar, BESS, and EV transactions across commercial, industrial, and public-sector clients in the U.S., UK, Canada, and India, including a 30MW solar project for a major tech company, a multi-site retail solar portfolio spanning national grocery and consumer brands, securitized solar financing structured with Deutsche Bank and Rabobank, distributed EV charging deployments across commercial campuses, and school district solar procurements underwritten under federal ITC frameworks.

His research draws directly from this deal experience, with a particular focus on cost benchmarking gaps, PPA pricing mechanics, and the policy variables that determine what projects actually get financed and built. He writes for practitioner outlets including Global Banking & Finance Review, Energetica, RealClear Energy, and E+E Leaders for Tomorrow.

Raj holds an M.S. from Duke University's Fuqua School of Business.

He can be reached at https://www.linkedin.com/in/kshitizraj/