Advances in Consumer Research
https://acr-journal.com/
Volume-2 | Issue-4 | September 2025

Original Researcher Article

E-Commerce vs. Traditional Retail: A Data-Driven Comparison of Profitability and Sustainability

Rahul Brahmbhatt

President, SSR Group, Arizona, USA
Email: barot81277@gmail.com

Received:10/08/2025
Revised: 20/08/2025
Accepted:12/09/2025
Published:30/09/2025

ABSTRACT

This work compares and contrasts unit economics and greenhouse-gas intensity of e-commerce, store-only, and omnichannel/BOPIS retailing over a multi-year panel (2018–2024) spanning North America and Western Europe through the merger of firm-year financials with order-level operations, geospatial logistics traces, and normalized emissions factors to construct bottom-up estimates of contribution margin, EBITDA, ROIC, and kgCO2e per order and per dollar of revenue. The empirical method combines lifecycle carbon accounting and activity-based cost-to-serve models with panel econometrics—specifically, staggered difference-in-difference with event-study diagnosis, instrumental-variables specifications, and dynamic panels—to measure causal channel outcomes and their durability. Decomposition techniques (Oaxaca–Blinder, Shapley) and causal mediation attribute shares of gaps among transport, facility energy, packaging, and returns; heterogeneity is described along dimensions of urban density, basket size, return propensity, fulfillment topology, and vehicle mode (ICE/EV/cargo bike). Operational experiments complement econometrics: vehicle routing simulation (CVRPTW) estimates cost and emissions under depot density, delivery window, and mode mix; field A/B testing investigates returns policy and pickup prompts. Managerially, the work provides break-even contours for basket size, route density, and carbon, as well as prescriptive levers for slot pricing, inventory placement (DC/micro-fulfillment/ship-from-store), packaging right-sizing, and last-mile electrification design. Collectively, the study offers an integrated, verifiable framework linking financial and environmental ledgers at order granularity, and enables channel decisions that expand the joint profit–planet frontier. Scope encompasses cradle-to-customer plus return, with precise boundaries and cluster-robust uncertainty reporting by cohorts regionally.

Keywords: E-commerce, Unit economics, Lifecycle carbon accounting, Greenhouse-gas intensity, Difference-in-differences.

INTRODUCTION

The demand moves across channels, and retailers are forced to protect margins and minimize environmental externalities. This work contrasts three archetypes: (i) pure e-commerce orders made online and delivered to the home through distribution or micro-fulfillment centers; (ii) store-only orders made online and delivered at the brick-and-mortar location through customer travel; and (iii) omnichannel/BOPIS orders made online and delivered at the store by picking it up or doing curbside. The unit of analysis is two-fold, namely retailer-year to represent firm performance and order to represent operational microeconomics. The location is North America and Western Europe, and the time is 2018-2024. Constructions are cradle-to-customer, and returns minus manufacturing emissions are shared in all channels. The operational scope incorporates picking, packing, facilities energy, line-haul, last-mile, payment processing, packaging, and reverse logistics. The non-parametric differences in categories are handled by employing fixed effects to adjust margins and return propensities.

The research measures net profitability channel differentials at the order and retailer-year levels. The leading performance indicators are contribution margin per order, EBITDA margin, and the return on invested

capital. Costs are assigned through activity-based drivers: seconds per pick and pack, kilometer line-haul and last-mile, payment fees per transaction, and anticipated write-downs due to returns. Meanwhile, greenhouse-gas intensity is calculated in kilograms of CO_2-equivalent per order and revenue and includes transport, facility energy, and packaging components. The analysis estimates the average treatment effects of the channel on these results. It describes heterogeneity in terms of urban density (number of households per square kilometer), basket size, and the rate of return. It also determines thresholds - such as delivery stops per hour or picking up bundling probabilities - at which a single channel becomes profit- or carbon-dominant, and indicates statistical uncertainty with cluster-robust standard errors.

This contribution to the literature and practice is threefold. It offers a joint profit-carbon lens, reporting results in both monetary and carbon emissions units, with uniform system boundaries, allowing for objective multi-optimization and avoiding siloed metrics. The study also employs plausible causal designs, including staggered difference-in-differences with event-study diagnostics and cohort-specific estimators, to capture treatment timing. These designs are supplemented by firm-clustered standard errors and placebo tests, which enhance the internal validity relative to the descriptive comparisons. The paper presents an operational decomposition that assigns observed differences to transport, facilities, packaging, and returns using Oaxaca-Blinder and Shapley decompositions; the findings are presented in unit economics and carbon waterfalls. Collectively, these contributions contribute to actionable advice on route density, pickup design, packaging right-sizing, and returns policy, and provide evidence on the extent to which urban form and energy mix affect the profit-planet frontier.

The channel choice changes the level and variance of cost-to-serve (CTS) and carbon-to-serve (C2S). Where D refers to the delivery area density, B to the basket size, and R to the return rate. In pure e-commerce, CTS (B, D, R) declines with D due to shorter kilometers per stop and increases with R due to reverse logistics and markdown risk. C2S also declines with D and a high electric-vehicle share, and increases with packaging mass. Store-only CTS relies on labor productivity and fixed-cost absorption, whereas C2S relies on customer round-trip distance and vehicle efficiency; multipurpose trips reduce emissions per order. In Omnichannel/BOPIS, CTS relies on picking productivity and pickup queuing, whereas C2S relies on the bundling of pickup with routine travel. Inventory placement and topology mediate each other, and such confounders as marketing intensity, product mix, and energy mix are involved.

This work is structured into various chapters. Chapter 2 is a summary of engineering evidence on retail unit economics and lifecycle assessment, with pitfalls including unobserved selection and inconsistent boundaries. Chapter 3 outlines data sources, integration, variable building, and quality control, which involve financial reconciliation, outlier treatment, imputation, and telemetry validation. Chapter 4 defines identification strategies and econometric specifications, such as event-study plots, staggered estimators, and decomposition techniques, that explain the gap in transport, facilities, packaging, and returns. Chapter 5 presents the results, including uncertainty intervals, heterogeneity in aspects of density, basket size, and rate of returns, as well as the break-even distances for carbon. Chapter 6 makes managerial inferences, scaling coefficients into route-density objectives, pickup design guidelines, package cuts, and returns interventions. Chapter 7 summarizes future work on real-time carbon and multi-echelon inventory, and Chapter wraps up the implementation checklists on omnichannel networks.

LITERATURE REVIEW

Retail Cost Structures & Unit Economics

The examination of e-commerce and traditional retail involves a data-driven comparison of costs, presented as an exact map of the expenses at the order level and how they are aggregated to inform the firm's performance. In store-based methods, contribution margin per order = price less COGS, store labor, shelving and checkout, shrink, payment fees, utilities, and an occupancy allocation [37]. In e-commerce, the same measure further subtracts pick-pack labor, packaging, line-haul and sortation, last-mile, anticipated return expenses, and customer-service handling. In abbreviated form: CM order = price [?] (COGS + pick/pack + packaging + line-haul + last-mile + payment fees + expected returns + CX). At this granularity, modeling allows channel-level waterfall analysis to assign profit erosion to specific operational drivers, including increased payment interchange, increased parcel weight, or decreased control density.

As shown in the figure below, the unit economics waterfall synchronizes store versus e-commerce channel order-level expenses and accounts for the contribution margin per order. The cost of store-based sales is reduced by prices minus COGS, store labor, shelving/checkout effort, shrink, payment fees, utilities, and occupancy allocation. The e-commerce bar also subtracts pick-pack labor, packaging, line-haul/sortation, last-mile delivery, expected return spend, and customer-service handling. This chart facilitates channel-level profit diagnostics, attributing margin erosion to operational drivers, including payment interchange, heavier-weighted parcels, and lower control density.

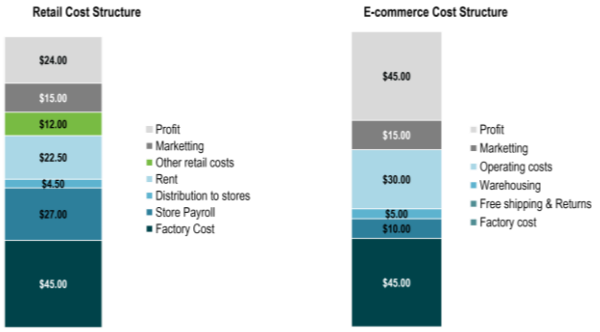

Figure 1: Unit-economics waterfall by channel: store vs. e-commerce

Scale economies manifest through different channels. Fixed rent and management overhead are diluted in stores by footfall, and significant fulfillment centers use wave planning, batch picking, goods-to-person automation, and sorter throughput. The two channels have the advantage of economies of purchase imposed by COGS [38]. However, e-commerce also faces diseconomies of low route density: the cost per stop in the last mile increases towards the outskirts of the territories and during periods of low demand when consolidation is lost. A realistic break-even analysis establishes contribution margin, store and e-commerce, on equal footing and problems out the minimum basket size or stop density needed to achieve store economics of delivery to equal a shopper trip, mode, service level, and probability of return. The managers can then model the impact of interventions such as slot pricing, micro-fulfillment, or dynamic bundling, on the threshold at which delivery is a win.

Real-world unit economics require the consistency of low-latency localized data that binds together transactional, operational, and telemetry information. Since order, warehouse, and transport management systems have complex nested payloads, retailers are storing order events, inventory mutations, and routing outcomes in document-oriented databases to avoid them being lost. Explicit write-read guarantees, versioning, and conflict resolution enable margin computation to seamlessly integrate asynchronous line-item states with payments and refunds, eliminating the need for manual counting. This functionality means in practice idempotent ingestion of events, persistent audit trails, and explicit consistency levels that are sensitive to the workload, all of which support faithful per-order profitability measurements [10].

EBITDA and ROIC are placed above the contribution margin and are disciplined in the allocations. EBITDA includes central G&A, technology OPEX, and facility overhead, which are not included in the order-level table. However, it does not include depreciation and amortization. ROIC has uniform capitalization rules over leasehold improvements, automation, vehicles, and software, and a rolling NOPAT numerator adjusted to returns, chargebacks, and breakage. A defensible pipeline balances the amount of order-level contribution margin to audited revenue and gross profit, and overlaid overheads to balance the general ledger, and the same design supports CLV modeling, balancing margin expansion to CAC, discount rate, and churn.

Sustainability of Retail Channels

The sustainability comparison depends on the boundaries of the lifecycle accounting and on the physics of movement and materials. Analysis should at least cover cradle-to-customer effects and a reverse logistics treatment of returns. Facility energy, transport, packaging, and returns can be used to break down per-order greenhouse-gas intensity. The facility energy allocates kWh to each order based on throughput and multiplies it by the local grid factor. Transport = (distance in legs) / load factor multiplied by a mode emission factor. Multiplication of material masses in packaging is through embodied factors. Returns combine anticipated emissions of reverse legs, restock processing, and disposal.

As shown in the figure below, the sustainability comparison is structured as a lifecycle accounting framework that calculates the per-order greenhouse-gas intensity across four elements: facility energy, transport, packaging, and returns. Facility energy accrues to each order through throughput (cubic throughput or orders) times the local grid emission factor. Transport emissions are computed per leg as the traveled distance divided by the load factor, multiplied by a mode-specific emission factor, and then summed over the route. Packaging impacts arise from multiplying each material's mass by its corresponding embodied carbon factor. Returns cover anticipated emissions from reverse-logistics legs, restocking and refurbishment processing, and end-of-life disposition or recycling impacts.

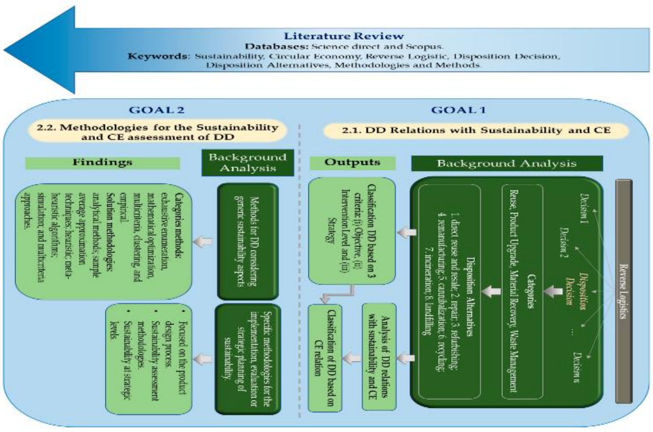

Figure 2: Retail lifecycle breakdown: facility energy, transport, packaging, and returns.

There is variation in operational reality that has to be modeled in event data and telemetry. Attempts to deliver, customer-not-at-home rates, re-dispatch, cross-dock transfers, weather-induced detours, buy-online-pick-up-in-store, and curbside shifts all increase emissions and cost; queueing dynamics are of interest in shopper trips and parking-lot idling. These micro-events are ideally monitored using event-driven data pipelines, which reduce latency, ensure exactly-once semantics with idempotency keys, back off on retries, and maintain durable logs to support consumer offsets. This enables trustful lifecycle accounting at operational granularity [6]. The evolution of the schema and watermarking can deal with out-of-order edge device arrivals, and the windowed aggregation can close the daily carbon ledgers without discarding late updates.

Special attention should be paid to the packaging and reverse logistics. Temperature-controlled or fragile items require either heavier or multi-material solutions that increase emissions and harbor risk, presenting a nonlinearity in both price and footprint. The computation of the expected return at order creation is a combination of category return probabilities and reverse routes and refurbishment paths modelled. In case refurbishment is not successful, it should consider the factors of recycling yield or landfills. Since policy-endogenous variables include the return propensity, any assessment must instrument or randomize policy-administered variables, such as window length, fees, and restocking policies, to prevent biased footprint estimates.

Improvement does not come by measurement alone. The parameterization of the routing optimizers must trade off distance, service levels, and load factor; the depot topology and location of the stock must be jointly optimized with delivery windows, and the EV, cargo-bike, or walking courier should be targeted to the neighborhood where density and curb-access constraints are favorable to high stop rates. HVAC schedules, lighting controls, and automation duty cycles need to be aligned on the facility side with order waves to minimise kWh per order [30]. These levers together are linked directly to the transport and facility parts of the lifecycle model and can be tested in controlled field pilots.

Causal Evidence in Channel Strategy

Descriptive comparisons of cost or emissions are also misleading when there are unobserved shocks that are accompanied by channel adoption. Causal identification is thus needed to separate the channel effects and confounders. Staggered difference-in-differences matches adopters of e-commerce, BOPIS, or micro-fulfillment to never-treated controls and allows the timing of treatment to vary. Event-study plots address pre-trends and display dynamics, such as short-run productivity dips during ramp-up and efficiency gains. Simple two-way fixed effects incorporating negative-weight artifacts can be circumvented by modern estimators that use cohort-time effects and cluster-robust standard errors, which account for the firm or region effect. These controls include the size of the basket, mix of categories, marketing intensity, weather, and fuel prices; unobservables that are constant within firms or periods are absorbed by firm and period fixed effects.

Complementary leverage is provided by natural experiments that are influenced by disruption. Unanticipated store-level closures due to extreme weather, infrastructure failure, or public-safety events generate plausibly exogenous channel availability, allowing the comparison of customer behavior, last-mile load, and profitability in impacted and unimpacted areas. Such designs need mature continuity management: structured runbooks, communication trees, escalation paths, and post-incident reviews are essential to allow the exposure windows, geographic scale, and counterfactuals to be recorded at a precise enough level to be inferential [28]. In addition, time stamps of outage initiation and recovery, rebalancing of inventory, and replacements are also recorded as part of disciplined incident logging, enabling researchers to match operational telemetry to monetary effects.

Credibility is focused on robustness. To guard against spurious results, analysts ought to re-estimate findings using alternative matching sets, geography-time cells, and exposure definitions, to conduct leave-one-cohort-out analyses, and to use randomization inference to guard against the presence of spurious results. Sensitivity analyses should be used to change the emission factors, the allocation of the facilities, and the model of the return probabilities, and to challenge the assumptions that conclusions rely on [40]. The additional checks are provided in the form of placebo interventions during pre-treatment periods and falsification tests on results that do not move. In cases where the risk of endogeneity exists, exogenous variation can be introduced through tools such as staggered rollouts of fiber, changes in delivery-zone boundaries, or carrier capacity limits, where relevance and exclusion restrictions are warranted and weak-IV diagnostics are reported.

Heterogeneity and mechanisms are essential for external validity. Quantile effects are used to illustrate the extent to which channel shifts alter the tails of profitability or emissions. Similarly, interaction terms are employed to determine the degree to which density, basket size, or return propensity mediate the results. Causal mediation analysis separates the impacts of treatment into those that can be ascribed to last-mile kilometers or packaging. Dynamic panels and distributed-lag models provide persistence and learning effects, and can be used together with state-space models to disentangle signal and measurement noise. False precision in managerial dashboards can also be avoided with variance-aware reporting, which includes confidence bands around event-study paths and prediction intervals around order-level footprints.

Gaps

Although measurement and identification toolkits have matured, there are still gaps at the interchange point of profitability, sustainability, and channel operations. There is little research on or systems that incorporate order-level financials and lifecycle emissions that are as fine-grained as those that are shared across identifiers. Consequently, finance and sustainability teams tend to draw conclusions based on various metrics, entities, and allocation bases [8]. An implemented solution includes a single event schema, whose order, item, facility, and vehicle keys are all nonmodifiable; product and customer attributes are all slowly changing; and reconciliation processes that ensure the carbon ledger is linked to the general ledger. This architecture allows the apples-to-apples allocation of the cost and footprint of the identical operational decisions.

Table 1: A Summary of Key gaps, impacts, and fixes

Area	Gap / Problem	Real-world Impact	Proposed fix / Implementation pattern
Measurement architecture	Profitability and lifecycle emissions are not integrated at order level; teams use different metrics and allocation bases.	Misaligned conclusions, inconsistent decisions, and inability to attribute profit/footprint to identical operational choices.	Single event schema with immutable keys (order, item, facility, vehicle); slowly changing product/customer dimensions; reconciliation linking carbon ledger to general ledger for apples-to-apples allocation.
Decision systems	Closed-loop, causal feedback rarely used in operations.	Pickers, customers, and fulfillment planners are not nudged toward low-cost/low-carbon choices.	Embed experiment-ready feedback: picker prompts for lighter dunnage, low-carbon delivery-slot offers when route density is high, store-led fulfillment triggers by neighborhood; design with fairness and spillover instrumentation.
Uncertainty handling	Data and model uncertainty do not propagate from collection to dashboards/A-B readouts.	Overconfident metrics; telemetry failure, imputed returns, and meter errors remain hidden.	End-to-end uncertainty accounting: quality flags, interval estimates, meter calibration error models, and propagation to KPIs and experiment analyses.
Optimization practice	Multi-objective optimization and Pareto frontiers are absent from everyday tools.	Managers lack transparent trade-off maps between contribution margin and $kgCO_2e$.	Integrate multi-objective solvers in routing/slot pricing; visualize Pareto fronts and iso-margin/iso-carbon curves for decision support.
Pricing signals	Carbon shadow pricing not consistently applied to routing or slot fees.	High-emission options appear cheaper; carbon and cost objectives diverge.	Introduce internal carbon price into objective functions and tariffs; expose marginal carbon costs in slot pricing and dispatch decisions.
Data standards & replicability	Inconsistent taxonomies for order state, transport legs, and packaging BOMs.	Weak reproducibility, limited benchmarking, and harder external audit/meta-analysis.	Adopt standard, open schemas; deterministic state machines; versioned taxonomies; reproducible pipelines with data hashes and audit trails.
Benchmarking & reporting	Lack of shared benchmarks and confidence-reporting conventions.	Slow learning cycles; risk of extrapolating local results to enterprise strategy.	Establish enterprise benchmarks, confidence/coverage standards, and out-of-sample validation rules; maintain metadata for study comparability.

Causal estimates of closed-loop feedback are hardly applied in decision-making. Adjacent areas have built tools that provide timely, structured, and individualized feedback and assess the behavioral and performance results using plausible counterfactuals. Applying that reasoning to retail implies systems where pickers are more likely to choose to select lighter dunnage, customers are offered delivery slots that are low-carbon when the route density is high, or store-led fulfillment is triggered in neighborhoods where it is the dominant choice on both cost and emissions. Importantly, this feedback must be experimentally or quasi-experimentally designed and instrumented in ways that reflect fairness and spillovers to be the best practice in developing and evaluating feedback interventions [16].

It has methodological blind spots that hamper adoption. Uncertainty does not tend to spread between data collection and dashboards and A/B tests; telemetry failure, imputed returns, and energy-meter error do not usually show up in reported periods. Multi-objective optimization and Pareto frontiers are not used in everyday tools and leave managers without a clear trade-off map between the contribution margin and kilograms of CO_2e [9]. Carbon shadow pricing is not consistently implemented in routing or slot pricing despite its direct relationship to both of these measures. Replicability has inconsistent order state, transport leg, and packaging bill of materials taxonomies; standard, open schemas would enable benchmarking, meta-analysis, and reliable external audit of channel strategies. Setting benchmarks and conventions of confidence reporting would speed up learning and minimise the risk of extrapolating local results to enterprise strategy.

METHODS AND TECHNIQUES

Data Sources & Integration

The interoperable datasets have been combined in this research to measure profits and sustainability at both the order and firm-year levels. The financial sources include the general ledger, trial balance, fixed-asset register, and order-level profit and loss allocations that assign COGS, picking and packing, line-haul, last-mile, payment fees, customer service, and return provisions to customer orders. The sources of operations are based on warehouse and order management systems (WMS/OMS): pick lists and timestamps, pickers' IDs, engineered labor standards, tote movements, parcel handoffs, transport management system (TMS) dispatches, and carrier invoices [44]. Facility energy meters (for electricity and thermal fuel) and any refrigerant logs (where applicable), packaging bill of materials with masses and dimensions of materials, and reverse logistics records (including reasons and dispositions of returns and secondary shipments) are all considered part of sustainability sources. Population density grids, road network graphs, land-use constraints, and weather histories are external sources and influence energy usage and delivery efficiency. Each source is timestamped and keyed so that order-level reconciliation and panel construction can be performed.

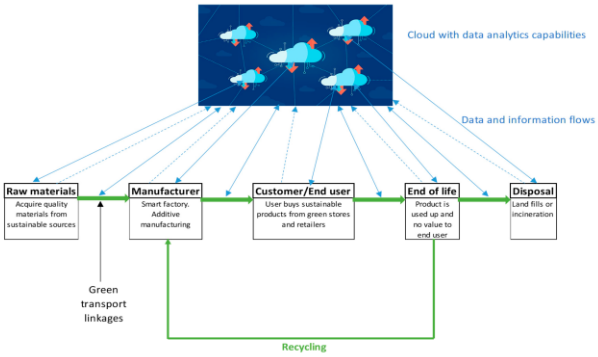

Figure 3: End-to-end supply chain data integration with cloud analytics

As shown in the figure above, the architecture collects data across the retailing lifecycle to measure profits and sustainability at order and firm-year levels. Financial data comprises the general ledger, trial balance, fixed assets register, and order-level P&L allocations for COGS, picking/packing, line-haul, last-mile, payment fees, customer service, and return provisions. Operation data are supplied by WMS/OMS (pick lists, timestamps, picker IDs, engineered standards, totes, parcel handoffs), plus TMS dispatches and carrier invoices. Sustainability data comprises facility electricity/thermal meters, refrigerant logs, packaging bills of materials, and reverse logistics records. External layers such as population density, road networks, land-use restrictions, and weather provide the context for efficiency. All are timestamped and join-keyed for audit and reconciliation purposes.

ELT consumes information in a controlled cloud warehouse. Order, orderlines, shipments, stops, returns, facility energy, vehicle energy, and payments are all canonical fact tables. Four dimensions are those of customers, items, facilities, vehicles, carriers, and geographies. Dimensions that change slowly (type 2) maintain changes in price, retrofit of facility, replacement of vehicles, and redefinition of carrier services [19]. Surrogate key-based referential integrity is implemented, and the business key (orderid, shipmentid, trackingid) is maintained to be reconciled with the operational systems. PII is tokenized, access is role-based, and column-level masked. Processing is coordinated around incremental snapshots, late-arriving data processing, and audit logs to achieve reproducibility.

Variable Construction & KPIs

The variables of profitability are built bottom-up. The contribution margin per order is defined as the selling price less the COGS, picking and packing, line-haul, last-mile, payment processing, anticipated returns, and customer service. The payment processing process is divided into ad valorem and per-transaction components, and observed incidence and recovery lags are used to amortize fraud losses and chargebacks. Expected returns cost is calculated by multiplying the predicted return by the scheduled transportation, handling, inspection, markdown, and disposal costs, minus the salvage value. The possibility of returns is calculated using historical cohorts in terms of product, customer, and promise-window characteristics. Cohort survival is used in calculating customer lifetime value (CLV) when the number of marketing customer acquisitions is amortized using cohort survival. The amortized number is customer marketing acquisition cost (CAC) x cohort survival = expected contribution margin x cumulative survival probability. Return on invested capital (ROIC) is calculated as net operating profit after tax divided by invested capital, where the denominator comprises net property, plant, and equipment (PP&E), right-of-use assets, and working capital related to the channel [21].

A modular bill-of-emissions assembles the sustainability variables at the order level. The sum of the Transport and Facility Share, Packaging and Returns corresponds to the total order greenhouse gas emissions. Leg mode aggregate emissions; at each leg, interface emissions are mode-specific emission factor times distance by realized load factor. Facility Share assigns site energy to orders through a causal driver, like completed picks, labor hours, or floor-time occupancy. Packaging emissions are aggregates of the multiplication of material masses by material factors. There are returns such as reverse-leg transport, reprocessing, repackaging, and scrappage.

Derived KPIs include kgCO2e per order, kgCO2e per revenue dollar, contribution margin per vehicle stop, and a carbon-adjusted contribution margin, which removes a shadow carbon price times order emissions. Break-even analyses find the solution to basket size, route density, or distance to customer pick up at which channel margins or emissions are equalized. Under mean assumptions, closed-form rearrangements are employed, and to obtain scale estimates, a Monte Carlo simulation is applied to model empirical distributions of basket size, propensity to return, and stop density [2].

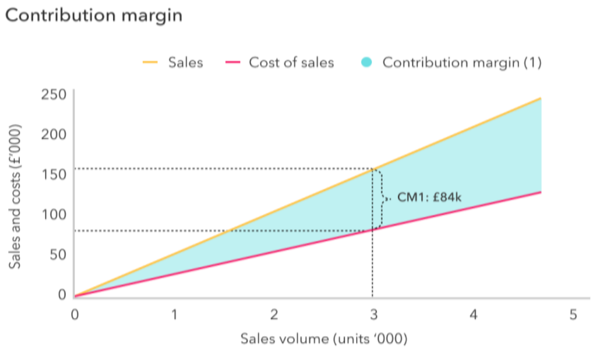

Figure 4: An overview of contribution margin and sustainability KPIs

As shown in the figure above, the light shade represents the contribution margin—the distance between sales and cost of sales—divided by sales volume. Contribution margin per order in this study is the selling price minus COGS, pick/pack, line-haul, last-mile, payment processing (ad valorem and per-transaction), foreseen returns, and customer-service handling. Foreseen returns cost combines foreseen return likelihood and transport, handling, inspection, markdown, and disposal, minus salvage. These order-level margins roll up to channel KPIs alongside ROIC and CLV (CAC amortized times cohort survival). The carbon-adjusted contribution margin subtracts a shadow carbon price multiplied by emissions. Break-even curves and Monte Carlo scenarios forecast basket size and route density.

Data Engineering

Engineering emphasizes reproducibility and inferential validity. Deterministic links between records include business keys, probabilistic when needed include fuzzy joins of customer name tokens, addresses, and timestamps, and clerical review of low-confidence matches. Units, currencies, and fiscal calendars are standardized; emission factors are time-indexed and versioned to capture updates to policy and grid. Verifying that the extremes represent a measurement error, not a rare occurrence, winsorized continuous cost, energy, and distance variables at the 1st and 99th percentiles [4]. Missing covariates are imputed by chained equations of multiple imputation, where the predictive mean matches in continuous fields and where the logistic or multinomial model is used in categorical fields; the imputations are performed through analysis to indicate parameter uncertainty.

Geocoding generates map-matched point locations that are matched to road graphs, and network shortest-path distances and travel times are calculated. Telemetry is smoothed to eliminate GPS drift and blasts of speed. Converting predictive analytics to operational reliability, Data pipelines are code, unit- and integration-tested, containerized, and deployed with continuous integration and delivery to maintain timely upgrades and reproducible experiments [23]. Stress tests of the scenarios are supplementary to the empirical data. Synthetic urban structures and depot structures are generated to explore edge behaviors in route density and accessibility. Generative models, trained initially on complex three-dimensional scenes, offer a principled means to sample plausible spatial configurations that emphasize routing choices and facility location [42].

Lifecycle & Boundary Choices

System boundaries assume cradle-to-customer and returns to match the managerial control. Direct (Scope 1) emissions are owned vehicle fuels and on-site combustion. Purchased electricity and steam (Scope 2) are aggregated on a facility scale with a temporal resolution, which should be good enough to capture seasonality and energy-mix variation [46]. Third-party carrier legs, packaging material production, leased facility electricity when sub-metered or contractually allocable, and outsourced fulfillment where billable energy or miles can be observed are relevant upstream Scope 3 elements.

The annualization of embodied emissions of durable assets (such as racking, vehicles) occurs as part of either vehicle electrification or replacement of racking decisions to material. Owned and contracted transport is explicitly partitioned to avoid double-counting, and facility totals are reconciled with carrier-billed energy or mileage to prevent discrepancies.

Validation & QA

Quality assurance encompasses both manual reconciliation and automated audits. Mass-balance tests require conservation: the number of orders placed is equal to the number of orders fulfilled and the number of cancellations and returns; the quantity of items carried out during pick, pack, ship, delivery, and reverse phases.

Financial reconciliation is a financial relationship between order-level costs and ledger totals by monthly account; residuals outside tolerance raise tickets and root-cause analysis. Carbon sanity checks compare scales of intensity to facility benchmarks. Monte Carlo draws on emission factors, load factors, and imputations propagate confidence intervals, accompany uncertainty, and point estimates [26].

THE CAUSAL INFERENCE AND PANEL ECONOMETRICS OF THE RETAIL CHANNEL EFFECTS.

Identification Strategy

This section describes the approach used by the study to isolate causal impacts of retail channel design on profitability and greenhouse-gas (GHG) intensity. The main design is staggered difference-in-differences in a multi-firm panel. There are e-commerce or buy-online-pickup-in-store (BOPIS) at other dates than some of the firms, and a subgroup of those that are store-only during the window [34]. Adoption is described as the quarter where at least ten percent of the orders are fulfilled online or through pickup, and this percentage remains consistent in the next two quarters. The estimand is the average treatment effect on the treated at the firm-market level, where markets are characterized by commuting zones to observe local person interaction methods are used to determine the weighted average of staggered treatments.

The natural experiments are complementary to quasi-random variation. The marginal contribution of physical outlets to omnichannel systems is determined by exogenous store closures that are not due to local demand, such as lease expiries without relocation. The rollout of low-emission zones changes last-mile mode choice and routing; border versus treated zones provide a policy discontinuity. Carrier fuel price shocks produce high-frequency price movements in delivery cost; exposure shares by carrier produce difference-in-exposure instruments [3]. Each shock may be coded in a DiD frame or serve as a channel intensity instrument. This is identified either by parallel trends in the untreated potential outcomes or by the exclusion restriction that the shock may only influence outcomes via the channel pathway and not directly. Automated checks, versioned data artifacts, and continuous testing are used to implement reproducibility and integrity of the estimation pipeline. The implementation of logistics environments. The identifying assumption is that, without adoption, both the contribution margin per order and $kgCO_2e$ per order would have moved in parallel between treated and control units.

The pre-trends are tested and dynamics are mapped using event-study estimators. Leads and lags against the adoption quarter trace coefficients of 8 quarters before and after adoption. The parallel trend assumption is supported by pre-treatment coefficients that are individually small and collectively negligible; the post-adoption profiles show ramp-up, plateau, or decline. The quarter of adoption is never counted to prevent the transients associated with launch marketing or a non-recurring capital expenditure. Event-study weights are made cohort-level and are aggregated using interaction-weighted aggregation to avoid composition bias, as in-

DevSecOps, such as static code analysis, dependency scan, and automated test gate, minimizes the configuration drift and bolsters the audit trail of causal claims [20].

Model Specifications

The baseline estimator is a two way fixed effects panel model: $y_{it} = \alpha_i + \lambda_t + \beta \cdot Ecom_{it} + X'_{it}\gamma + \varepsilon_{it}$, where i indexes firm or firm market and t indexes time, as highlighted in Table 2 below. Ecomit is a measure of treatment or an intensity measure, like the percentage of online orders or the number of BOPIS. Xit has assortment mix controls, marketing spend controls, seasonality controls, fuel price controls, weather controls, and macro conditions controls. The interaction-weighted estimators of staggered adoption, to recover cohort-specific and dynamic effects, prevent negative weighting and obtain interpretable group-time average treatment effects. Inference applies firm-level cluster-robust standard errors; in cases where the number of clusters is limited, wild cluster bootstrap p-values are provided.

Table 2: Econometric model specifications, controls, and diagnostics

Approach	Specification / Key Equation	Key Controls / Instruments / Assumptions	Diagnostics, Inference & Outputs
Two-way fixed effects (baseline)	$y_{it} = \alpha_i + \lambda_t + \beta Ecom_{it} + X'_{it}\gamma + \varepsilon_{it}$; *i* i=firm (or firm–market), *t* t=time. Ecom =treatment/intensity (e.g., % online orders, # BOPIS).	Xit assortment mix, marketing spend, seasonality, fuel prices, weather, macro conditions. Identification: parallel trends conditional on FE and Xit.	Firm-level cluster-robust SEs; wild-cluster bootstrap p-values when clusters are few. Outputs: average treatment effect β and fitted counterfactual paths.
Interaction-weighted staggered adoption	Group-time ATTs via interaction-weighted/event-study estimators yielding dynamic β_k (leads/lags) by cohort; avoids negative weights from naïve TWFE.	Same covariates as baseline; treatment timing varies across cohorts. Assumes no anticipation and cohort-specific parallel pre-trends.	Pre-trend tests on leads; cohort-specific effect paths $k \in [-8,+8]$; stacked designs with never-/later-treated controls; visualization of $\hat{\beta}_k$.
Dynamic panel GMM (Arellano–Bond/Bover)	$y_{it} = \rho y_{i,t-1} + \beta Ecom_{it} + X'_{it}\gamma + \alpha_i + \lambda_t + u_{it}$. Endogenous $y_{i,t-1}$ instrumented with deeper lags	Handles persistence from learning, demand maturation, fleet transition. Monetary outcomes deflated to constant	Hansen J for over-ID; AR(2) serial-correlation test; report long-run multiplier $\beta/(1-\rho)$ with delta-method CIs. Sensitivity to

Approach	Specification / Key Equation	Key Controls / Instruments / Assumptions	Diagnostics, Inference & Outputs
	(difference/system GMM).	currency; winsorize 1st/99th pct.; carbon intensities normalized per order and per revenue dollar.	instrument count/finite-sample bias.
Instrumental variables (2SLS/IV)	First stage: $\mathrm{Ecom}_{it} = Z_{it}{}'\pi + W_{it}{}'\theta + \eta_{it}$Ecomit=Zit′π+Wit′θ+ηit. Second stage: $y_{it} = \alpha_i + \lambda_t + \beta\,\widehat{\mathrm{Ecom}}_{it} + X_{it}{}'\gamma + \varepsilon_{it}$.	Instruments Zit: historical broadband coverage, distance to new fulfillment centers, exposure to logistics-relevant policies; controls Wit complement Xit. Assumes instrument relevance and exclusion.	Relevance via Kleibergen–Paap rk F; exogeneity via Sargan–Hansen; weak-IV-robust Anderson–Rubin and conditional LR inference when first-stage is marginal. Outputs: βIV with robust CIs.

Dynamic panels exhibit persistence in profitability and carbon intensity, which are driven by learning, demand maturation, and fleet transition. The generalized method of moments models developed by Arellano, Bond, and Bover also have an additional lagged dependent variable and have instrument internal dynamics with deeper lags. Diagnostics such as the Hansen J-test to over-identify restrictions and the AR(2) tests to check serial correlation of the residuals. Dynamic specifications enable estimates of long run multipliers $\beta/(1-\rho)$, where ρ is the autoregressive coefficient [22]. Monetary results are deflated to constant money and winsorized on the first and ninety-ninth percentiles; the carbon intensities are normalized by order and revenue dollar to ensure cross-firm and cross-period comparability.

The instrumental-variables estimation directs simultaneity between channel intensity and outcomes. The former step models Ecomit as a function of instruments Zit - historical one broadband coverage, distance to newly opening fulfillment centers, or varying exposure to logistics-relevant policies - and is controlled by Wit. The second step replaces fitted values in the outcome equation to estimate β. Relevance is tested based on Kleibergen-Paap rk F-statistics; exogeneity is tested based on Sargan-Hansen tests [47]. When first-stage strength is marginal, weak IV-robust tests, such as Anderson-Rubin tests and conditional likelihood ratio inference, are reported.

Decomposition & Mediation

Structural accounting is combined with statistical decomposition, attributing observed gaps to operational drivers. A cost per-order waterfall is built: cost of goods sold, line-haul, picking and packing, last-mile, customer service, and expected returns. An equal footprint waterfall assigns transport emissions, facility energy, packaging, and returns. Oaxaca-Blinder decomposition is a measure of the proportion of the mean channel gap that is attributable to differences in these observed elements, compared to unexplained residuals. Invariance to ordering Shapley value decompositions constitutes a robustness check that is invariant to ordering.

Isolation of mechanisms: Causal mediation analysis. Noted last-mile kilometers per order and the return rate. Mit can be the set of candidate mediators. Natural direct and indirect effects are obtained by estimating the treatment-to-mediator and mediator-to-outcome relationships with the proper controls [24]. In the case of probable mediator endogeneity, they employ an IV-mediation design: topology or policy shocks, including the opening of depots, a change in delivery-slot pricing, or tightening of return-windows, are instruments of M_it. A sensitivity analysis is an unobserved mediator-outcome confounding that bounds the indirect effect.

Robustness & Sensitivity

A large falsification battery supports the credibility. To test spurious pre-effects, cohorts that have not yet been treated are placed in a placebo. Sensitivity is checked by alternate exposure thresholds, such as adoption at five or twenty percent online share. The leave-one-firm-out and leave-one-market-out re-estimations address the concentration risk. Oster's δ quantifies the strength of unobservables required to explain away β given movements in R^2 and coefficients as controls are added; values exceeding unity indicate that unobserved selection would need to be implausibly strong to overturn conclusions [12]. Permutation tests implemented in randomization inference do not destructure the panel. The influence diagnostics roll away extreme return-rate deciles and re-fit all models.

Since event-studies, bootstrap, and grid-IV-GMM are computationally intensive, the analysis implements cost and scale guardrails similar to those in microservices governance, including job budgets, autoscaling, and caching of adequate statistics to prevent recomputation. This discipline is comparable to best practices in distributed systems [7].

Heterogeneity & External Validity

Policy relevance needs to align with the areas where e-commerce is prevalent. Heterogeneity is estimated in terms of interactions with density, depot topology, basket size, and temperature, mode of delivery, and customer distance and quantile treatment effects. External validity is tested by re-estimating on hold-out geographies and by reweighting to target markets.

Collectively, these components provide decision-quality estimates of actual retail networks and scalable policy simulations in systemic form, in the country.

EXPERIMENTS AND RESULTS

Experimental Designs

The program interweaves quasi-experimental econometrics and operational experiments to extract channel-specific profits and greenhouse gas impacts apart from the base channel impacts. An event study is centered on the quarter of initial e-commerce delivery or purchase-online-pick-up-in-store (BOPIS) initiation. Firm-quarter results for every treated firm are cohort-aligned and stacked with never-treated or subsequently-treated controls. As presented in Table 3 below, dynamic coefficients β_k are calculated for $k = -8…+8$ by firm and time fixed effects, cohort dummies, and seasonal and regional macro covariates. Diagnostic checks for parallel trends are performed with pre-period coefficients and visual checks; leads/lags are corrected via winsorization to prevent leverage of extremes. The primary outcomes are contribution margin per order, return on invested capital (ROIC), purchase return rate, and kg CO_2e per order [25].

Experiment	Design & Method	Key Levers / Controls / Constraints	Primary Outputs & Metrics
Event-study around channel launch	Quasi-experimental event-study centered on the quarter of initial e-commerce delivery or BOPIS initiation. Firm–quarter panels are cohort-aligned and stacked with never-treated or later-treated controls. Dynamic coefficients β_k are estimated for $k \in [-8, +8]$ using two-way fixed effects with cohort dummies and seasonal/regional macro covariates. Parallel-trend diagnostics use pre-period coefficients and visual checks; leads/lags are winsorized to reduce extreme leverage.	Fixed effects at firm and time; cohort indicators; seasonality and regional macro controls; winsorization of leads/lags; diagnostic checks for pre-trends and specification stability.	Contribution margin per order; ROIC; purchase return rate; kg CO_2e per order.
Routing scenarios (CVRPTW)	Operational simulation of a capacitated vehicle-routing problem with time windows. Geospatial demand is seeded from historical orders; congestion captured via time-dependent speeds. A large-neighborhood search (ruin-and-recreate with adaptive penalties) constructs near-optimal tours.	Scenario levers: depot density (DC only; DC + micro-fulfillment; ship-from-store); vehicle mix (ICE van, EV van, cargo bike); delivery windows (all-day, 4-hour, 1-hour). Constraints: vehicle capacity, shift length, service time. Emissions via mode-specific factors and load-allocation rules.	Cost per stop; stops per route; vehicle and labor utilization; lateness/service-level attainment; route-level kg CO_2e.
Returns-policy A/B test	Cluster-randomized controlled trial at the store-cluster level to prevent spillovers. Two arms: lenient (free postal returns, 45-day window) vs. strict (in-store-only, 14-day window, tiered restocking fees). Power analysis targets a 2 pp minimum detectable effect on return rate and 5% on restock cost.	Randomization at cluster level; treatment integrity monitoring; governance controls: least-privilege access, rigorous authentication, and audit trails to safeguard data integrity and inference validity.	Experienced return rate; restock cost per returned unit; additional transport kilometres; processing lead time; incremental reverse-logistics emissions.

Routing scenarios describe the operational frontier with a capacitated vehicle-routing problem with time windows (CVRPTW). Scenario levers are depot density (distribution center only; DC + micro-fulfillment; ship-from-store), vehicle mix (internal-combustion van, electric van, cargo bike), and promised delivery windows (all-day, four-hour, one-hour). Seeding geospatial demand is based on historical orders, while congestion is captured through time-dependent speeds. A large-neighborhood search with ruin-and-recreate and adaptive penalties yields near-optimal tours under vehicle capacity, shift length, and service-time constraints. The simulator for each scenario yields the cost per stop, stops per route, utilization, lateness, and route-level kg CO_2e emissions under mode-specific emission factors and load-allocation rules.

A returns-policy A/B test estimates causal effects of lenient versus strict rules. Randomization happens at the store-cluster level to avoid spillovers. The lenient arm offers free postal returns and a 45-day window; the strict arm limits returns to in-store returns with a 14-day window and tiered restocking fees. Primary outcomes are experienced return rate, restock cost per returned unit, additional transport kilometres, lead time for processing, and incremental emissions for reverse logistics [13]. Power analysis centers on a minimal detectable effect of two percentage points for the return rate and five percent for the restock cost. Experiments all apply least-privilege access, rigorous authentication, and audit trails for safeguarding data integrity and upholding inference validity [29].

Profitability Results

Results are presented as coefficients, waterfalls, and break-even charts. The table reports intent-to-treat effects from an event study and average treatment effects from a staggered difference-in-differences analysis. As highlighted in Table 4 below, column (1) includes firm and time controls; Column (2) also includes category mix and promotions; Column (3) also includes marketing spend and regional unemployment; Column (4) uses exogenous broadband rollout and distance to newly created fulfillment nodes as an instrument for channel penetration. The models report cluster-robust intervals and standardized effect sizes. Post-estimation converts the coefficients into annualized deltas of profit at observed volumes and into movements per order unit economics.

Table 1: Profitability Results Reporting Plan—Models, Inputs, and Outputs

Module	What it covers	Inputs & Methods	Key Outputs
Econometrics (Cols 1–3)	Intent-to-treat event-study and staggered DiD with firm/time fixed effects; adds category mix, promotions, marketing spend, regional unemployment	Quarterly panel; outcomes: contribution margin/order, ROIC; lagged controls; cluster-robust SEs; standardized effect sizes	β estimates with 95% CIs; partial R^2 changes; stability of signs/magnitudes
IV Estimation (Col 4)	Instruments channel penetration to address endogeneity	2SLS using exogenous broadband rollout and distance to new fulfillment nodes; first-stage F; over-ID tests	LATE β; weak-IV diagnostics; robustness vs. alternative IV sets
Post-Estimation Translation	Converts coefficients into business terms	Multiply β by observed volumes; map to per-order deltas; seasonal annualization	Annual profit delta (currency); per-order margin change with CIs
Waterfalls & P&L Reconciliation	Channel unit-economics decomposition and tie-out	ABC drivers (touches, minutes, throughput, floor area); allocate COGS, pick/pack, line-haul, last-mile, payments, CX, expected returns; trial-balance checks	Waterfall charts by channel; % variance vs. audited P&L within materiality thresholds
Break-Even & Mode Overlays	Minimum basket size/route density to match store-trip economics; mode comparisons	Iso-margin contours over stops/route × items/order with overlays for ICE van, EV van, cargo bike; BOPIS and ship-from-store adjustments	Channel/mode break-even thresholds; feasible regions and indifference zones
Operational Simulations	Micro-foundations for costs and constraints	CVRPTW with lateness penalties; elastic slot pricing; EV state-of-charge and depot power caps; cargo-bike payload/range; hybrid fleet assignment	Cost/stop and lateness trade-offs; density gains from demand shaping; EV feasibility schedules; hybrid allocation rules
Sensitivity & Robustness	Parameter uncertainty and identification checks	Tornado/probabilistic sweeps over wages, fuel/electricity, last-mile productivity; placebo launch dates; alternative FE estimators; ABC driver variants	Bands around iso-margins; no pre-trends; consistent signs/magnitudes; stable waterfalls under driver changes
ROIC & Bridges	Channel-level capital efficiency and sources of change	ROIC = NOPAT / invested capital (inventory, receivables, automation, vehicles, charging, platforms); working-capital turns; asset-productivity metrics	ROIC by channel; ΔROIC bridge into WC and fixed-asset contributions vs. hurdle rate

Contribution-margin waterfalls deconstruct store-only, delivery, and BOPIS into price, COGS, pick/pack, line-haul, last-mile, payment processing, customer care, and expected returns. Activity-based cost allocation allocates the common labor and facilities overhead by drivers (order touches, cubic throughput, handling minutes, and floor area). A reconciliation table reveals the quantities modeled, balancing the P&L within materiality thresholds through the audited P&L [5].

Break-even charts quantify the least basket size or route density wherein delivery serves store-trip economics. Iso-margin contours are plotted within average stops per route and items per order, with overlays for mode (ICE van, EV van, cargo bike). BOPIS contours add shopper travel distance, staging labor, and curbside dwell; ship-from-store contours add stockout externality and inter-store replenishment freight. Wage inflation, fuel/electricity prices, and last-mile productivity are depicted in sensitivity bands.

Operations simulations provide micro-foundations. Narrow delivery windows increase fleet costs and overtime, rising cost per stop and lateness risk; density or elastic slot prices offset the effect by shifting the demand. Electric vans reduce the variability of operating expenses but require schedules that respect charging dwell and depot power constraints; state-of-charge feasibilities determine routing construction. Cargo bikes are optimal within busy two-kilometer catchments but payload-limited; hybrids allocate heavyweight or long-haul portions to bikes and vans, respectively. To promote use, the reporting layer extracts post-hoc feature attributions for scenario meta-models and offers interpretable narratives tracing cost drivers back to levers [43]. Finally, ROIC is computed as after-tax operating profit divided by invested channel capital attributable to each channel (inventory, receivables, automation, vehicles, charging, and platforms amortized over useful life); delta-ROIC is linked to components from working-capital turns and fixed-asset productivity.

Sustainability Results

Environmental outcomes are illustrated through distributional graphs, component attributions, and break-even distances for carbon. The distribution of channel-wise $kgCO_2e$ per order, stratified by basket size and urban density, is illustrated using kernel density and violin graphs; tail risk through sparse routing, cold chaining, and oversized packaging is highlighted by dispersion. Uncertainty is propagated with confidence bands from metering of facilities, variability of emission factors, and multi-stop allocation rules.

Component attribution breaks order-level emissions into transport, facility share, packaging, and returns. Transport is equivalent to route kilometres multiplied by mode-specific factors, derived from stop service and capacity residue [1]. Facility share assigns kWh off of orders through throughput or time drivers selected by out-of-sample fit of meter readings. Packaging is equivalent to bill-of-materials mass time's material factors, minus damage-rate penalties if down-gauging material raises breakage. Returns the sum of the outbound and reverse legs plus re-processing energy, and attributes salvage value where refurbishment replaces production.

Carbon break-even distances make the channel decision operationally viable. The break-even point equates the courier's marginal routing emissions with a shopper's round-trip vehicle travel emissions, considering occupancy and alternative errands in the form of trip chaining. Hourly grid-intensity profiles for EV fleets induce time-of-day dependency, and hence bands are reported for grid cases. Time-of-delivery policies that promote consolidated delivery days, pickup-point aggregation, and locker networks encourage the break-even toward lower-carbon channels for low-density neighborhoods.

Heterogeneity & Mediation Findings

Heterogeneity is defined by basket size, density, temperature class, and return inclination. Quantile treatment effect results pinpoint the profitability premium of e-commerce as focused within large-basket deciles and high-density corridors, and store-only sustains the lead within exurban geographies with extensive deadheading. Complementary interactions for vehicle mix pinpoint that cost and carbon savings are reaped from electrification within medium-density locations, particularly where accompanied by micro-fulfillment or locking boxes. In sustainability, larger baskets and clumped stops reduce delivery $kgCO_2e$ per order, and high return inclination reduces both profit and footprint [18].

As shown in the figure below, the expenditure on products' Lorenz curve illustrates intense concentration: a small number of customers (the right tail) generate a large quantity of sales, which are concentrated in the large-basket deciles. This variability accounts for why the profitability benefit of e-commerce arises primarily in high-density corridors for high spenders, and store-only endures in exurban markets with a significant amount of deadheading. The same focus engages functional decisions: electric fleets, micro-fulfillment, and lockers generate disproportionate cost-carbon savings in middle-density locations; bunched stops and larger baskets reduce $kgCO_2e$ per order. Mediation analysis attributes variation to distance on the last mile and slope to returns, and returns eat into profit and footprint.

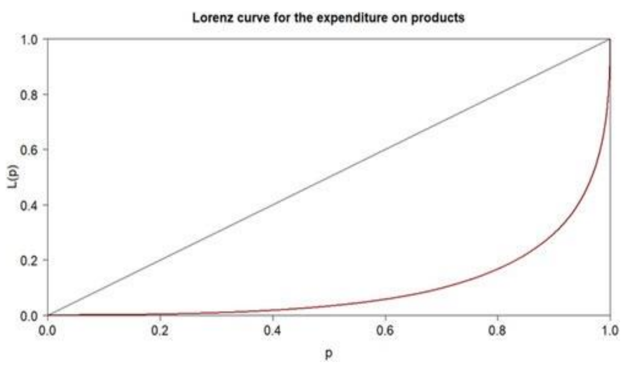

Figure 5: An Overview of Lorenz curve: large-basket deciles dominate spend in high-density areas

Causal mediation decentralizes channel effects into routes through last-mile kilometers per order and return rate. Sequential regressions with firm fixed effects and interaction of mediator and treatment distinguish natural indirect effects, quantifying the share of the margin gap and the carbon gap explained by preventable distance and policy-driven returns. Delivery days by zonal category, proper packaging, and different return windows by class are the top-affected levers, ranked by mediated effect.

Robustness Summary

Robustness checks answer identifications, measurement, and optimization assumptions. Pseudo-launch date placebo tests disclose no anticipatory margin or emissions trends. The other two-way fixed-effects estimators, which control for negative weighting under staggered adoption, produce similar signs and magnitudes. Randomization inference reaffirms the cluster-randomized returns A/B under spillovers over spaces [27]. Influence-function checks and extremal deciles of delivery distance and return rate trimming reassure that findings are not the product of outliers.

Measurement audits ensure that order-level carbon and margin amounts align with external standards and financial records. Utility bills are compared with facility energy meters; speed–distance abnormality verifications and ad-hoc spot time-and-motion studies cross-check telematics; vendor bills reconcile bills-of-materials for packaging. A reproducibility package—versioned code, frozen parameters, data hashes, and signed results—permits replication by an external party and prevents analytic drift upon re-run.

DISCUSSION

Managerial Implications

Data-driven lens puts channel decisions into concrete levers that can be ranked, priced, and tracked. Route density is the top motivator of last-mile cost and carbon. Managers need to pay attention to stops per route-hour, grams CO_2e per stop, and consolidation percent; set floor thresholds by market; and gate free delivery into zones off thresholds. Density can be engineered through synchronizing order cutoffs, coordinating delivery days, and preferring ship-from-store for short stems while protecting on-shelf availability. Density engineering also involves designing substitutions that maintain basket value while increasing the co-location of goods, thereby improving batching efficiency during execution. These rituals should be enshrined within weekly playbooks owned jointly by operations and merchandising organizations.

Micro-fulfillment, when deployed close to demand, can accelerate pick minutes per unit and shorten stem distance. However, the capital cycle needs to be warranted by a cost-to-serve model incorporating labor, depreciation, utilities, software, and cannibalization [15]. A phased rollout—pilot, densification, automation—mitigates risk. Payment fees are often disregarded. Calculate payment take rates by method and apply differential nudges, like small surcharges for high-fee wallets and rewards for low-fee bank transfers, and track authorization rates and churn. Smart slot pricing synchronizes customer adaptability and fleet utilization: charge a slot by incremental last-mile cost minus a consolidation credit based on forecasted added stops, and surge multipliers top off at protecting fairness and conversion.

Few strategies cut return rate faster than contributing and emissions uplift simultaneously. Start with a diagnostic that deconstructs returns into wrong item, damage, fit, remorse, and service failure. Use focused fixes for every cause: fit guides and size prediction, reinforcement of packaging for high-damage nodes, photo verification before drop-off, and exception playbooks. Test-and-learn cadences iterate through weekly data and monthly retrospectives. Upward improvements call for feedback loops involving input from stakeholders such that interventions remain relevant as the situation itself changes, a general design principle emphasized in feedback-centered program models [17].

Carbon levers are counterparts of profit levers, but they require measurable and auditable quantities of carbon. Mode shift for electric vans or cargo bikes should be considered marginal in terms of kg CO_2e/km by vehicle and grid intensity, and reoptimized monthly as the grid and density vary [45]. Payback occurs for electrification quickest where and when stop frequency and regenerative braking maximize circuit density. Material mass and dimensional weight are reduced by right-sized packaging. Managers ought to enshrine a cube- and drop-height-tolerance indexed fragility matrix for packaging selection and conduct controlled studies measuring damage versus material where the latter is the variable of interest. Ship-from-store is preferable where local stems outweigh linehaul plus delivery; otherwise, ship-from-DC prevails where it is feasible to consolidate lanes intensely.

Channel Strategy

Omnichannel models, chief among them buy-online-pick-up-in-store and curbside, are particularly successful in high-density catchments with numerous locations and short shopper trips. Similarly, pure delivery dominates where baskets are large and trips are much consolidated. A practical playbook diverges markets by urban core, near-urban, and exurban or rural markets [31]. Compute four ratios for each: store coverage (population share within a two-kilometer radius of a store), average basket value by channel, return propensity, and drivable consolidation opportunity (orders per square kilometer per delivery window). In places where the product of coverage and consolidation exceeds some policy threshold, prefer BOPIS and curbside; where basket and consolidation are the winners, choose home delivery.

Figure 6: Omnichannel retail flows: BOPIS, delivery, and store pickup

As shown in Figure 7 above, the omnichannel retail network synchronizes digital and traditional stores to serve three customer groups including digital shopping, buy-online-pick-up-in-store (BOPIS), and conventional in-store purchasing. For digital buyers, logistics assistance includes home direct delivery, while for BOPIS buyers, they purchase and pick up in locations, and traditional buyers make and pick up purchases at locations. Cost divides flow between locations and retail outlets, and capital and logistics flows between consumers and service means. This arrangement accommodates channel strategy direction: high-populated catchments benefit from BOPIS and curbside, and high-consolidation, large-basket markets benefit from pure delivery.

Inventory positioning must co-evolve with channel configuration. A two-echelon rule works: distribution centers hold depth for long-tail stock keeping units and promotion bulks, and stores or micro-fulfillment centers hold breadth for fast movers matched with local elasticity. Safety stock must factor in returns, lead times, and resale opportunities for waste reduction. Fulfillment rules require latency sensitivity: pick-up orders buffer stock upfront for cancellability, and delivery orders bunch together for five to ten minutes for more batching without losing service level [14]. Order orchestration must solve a constrained shortest-path with weighing of stockouts, stems, promised windows, and expected return handling.

Trade-offs and Co-benefits

Levers rarely move in isolation. Electricating the last mile reduces emissions but can increase route extension because of charge logistics, cold-weather buffer reserves, or charger queuing. Mitigate by routing with state-of-charge limitations, scheduling mid-shift top-ups at facilities, ordering eleven- to twenty-two-kilowatt chargers to facilitate thirty- to sixty-minute turns, and off-peak-time shifting of charge rates. Packaging minimization reduces material usage and volumetric weight, but can increase damage incidence and re-delivery miles [41]. A decision model should impose a damage-rate penalty on the unit economics waterfall such that packaging changes occur only when anticipated savings outweigh anticipated loss.

Returns policy overlaps with merchandising and customer experience. Narrow windows reduce the cost of handling and transport emissions, but may suppress conversion for commodities with uncertainty over their fit. The operational trade-off is to differentiate policies by predictability: big windows for ambiguous, high-margin, high-uncertainty goods and tighter terms for commoditized, fragile goods. Information flow is an implicit coupling across all these levers. Scalable, fault-tolerant comms systems compensate for coordination failure at peaks and incidents. Design perspectives on scalable comms emphasize feedback, reliability, and security where information crosses roles and settings; the same rules apply to connecting stores, depot, drivers, and service teams through comms-driven alerts, fault-tolerant audit trails, clear escalation channels, and few service targets [39].

Limitations

Some limitations define inference and implementation. Error in third-party linehaul measurements is omnipresent. Carriers can report zone-miles or linehaul miles instead of the actual length of the path, and multi-tenant linehauls make load-factor identification difficult. Enclose these inaccuracies by triangulating telematics subsamples with manifests and attributing uncertainty intervals to estimates of emissions. Discretionary unobserved shocks, such as bad weather, specials, and rival movements, can distort channel contrasts. Difference-in-differences with event windows and cluster standard errors is a good starting point. However, managers should also consider pretrends, perform placebo tests, and conduct leave-one-firm-out analyses for robustness checks.

Another constraint is the completeness of data. Packaging bills of material are typically absent for legacy stock-keeping units. Facility energy can be metered at the building level rather than the process level, and carriers' emissions factors can be generic rather than lane-specific. Impute responsibly, disclose the imprecision, and prefer conservative assumptions when you report externally. Generalizability also requires caution. Findings observed in dense, temperate cities do not automatically translate to sparse, hot locations where cooling loads and long stems are significant. Firms need to do market-specific pilots, report conditionally treated results segmented by density, basket size, and return frequency, and establish governance that attaches equal importance to both contribution margin and emissions per order such that single-metric gaming is deterred.

FUTURE WORK

Enriched Telemetry and Real-Time Carbon

Future projects should employ high-definition telemetry and facility metering to estimate marginal emissions at the point of service. Vehicle CAN bus streams, GNSS traces at up to 1 Hz, and warehouse IoT meters for HVAC, conveyors, refrigeration, and chargers can be combined to calculate energy utilization on a second-by-second basis. A successful streaming stack will include edge ingestion on vehicles and buildings, a durable message queue, and a document database for storage of semi-structured sensor payloads along with order identifiers. This enables online lifecycle accounting that multiplies instantaneous energy by time-varying emission factors (grid marginal intensity for power, fuel carbon intensity for Liquid Fuels), and allocates a proportionate share of facility energy for every pick and pack [35].

Data quality will be ensured through schema validation and idempotent writes; privacy will be maintained through the use of hashed customer identifiers, and geospatial precision will be constrained. Operationally, the platform should provide real-time key performance indicators—kgCO2e per stop, per order, and per basket value—and a shadow cost which lays bare the cost of an additional kg of CO2e in routing and slot-assignment services. To preserve scalability and retrieval for dashboarding and optimization services, the document database's sharding, along with time-series indexing, will be tuned for write-heavy ingestion and query patterns typical of telemetry workloads [11].

Multi-Echelon Inventory & Repositioning

Footprint and profitability both depend upon where the inventory is and where it moves. A multi-echelon optimization should co-design the distribution centers, micro-fulfillment centers, and stores, along with the carrier arrangements and fleet mixture. Decision variables will include safety stocks by SKU–location, cross-dock transfers, lateral transshipments, mode selection, and time-window servicing levels [32]. The objective will maximize the expected contribution margin minus lambda times emissions, where lambda is an internal carbon pricing corresponding to the firm's hurdle rate. Constraints will include region-wide servicing level, capacity within nodes, fleet and labor availability, and perishable windows.

The formulation can be a mixed-integer linear program with Benders' decomposition, deployed on a rolling horizon with weekly replans and day-by-day micro-adjustments, based on updated forecasts and realized actual demand. Repositioning policies need to be stressed with stochastic scenarios of fuel prices, return rates, and congestion; sensitivity analysis will provide value- of-relaxation for each constraint and dual prices for carbon. Interface with last-mile routing is via iterated feedback: the inventory model provides promised availability and cut-off times, and the vehicle routing problem provides marginal delivery cost and emissions by zone; fixed-point iteration continues until both are converged within tolerance.

Causal ML for Targeting

Heterogeneous effects must be measured to target interventions where they create value. Conditional average treatment effects for packaging right-sizing, return-fee thresholds, BOPIS prompts, time-window offers, and low-carbon delivery modes will be estimated by the program. A gradient-boosted outcome model and doubly robust corrections and policy learning under uplift losses will be stacked together via a meta-learner stack. As shown in the figure below, the causal ML workflow relates covariates to treatment and outcome. It distinguishes the averaged effect, effect heterogeneity, and individualized effect for binary and continuous treatments. The binary treatments for our research include BOPIS prompts, time-window offers, and low-carbon delivery nudges. Continuous "dose" choices are return-fee thresholds and right-sizing of packaging. Conditional average treatment effects (CATEs) are approximated by a meta-learner stack that combines gradient-boosted outcome models with doubly robust corrections and uplift policy learning. This approach involves cross-fitting, calibration, and out-of-fold uplift gain curves to validate the models. Released policies target customers, orders, and routes that are predicted to deliver the maximum profit and carbon impact at scale.

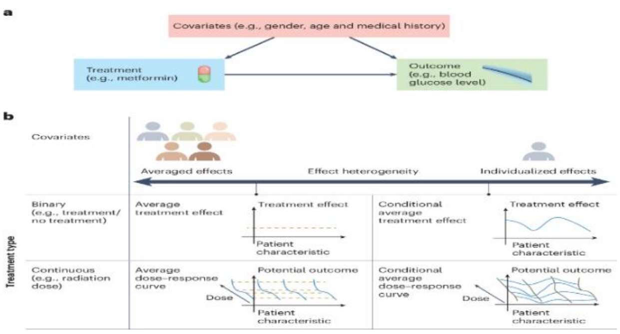

Figure 7: An Overview of Causal ML: estimating heterogeneous and individualized treatment effects for targeting

To maintain long-term behaviors through the seasons and customer lifecycles, the stack will consist of memory-based neural units writing and reading state for

past interactions; the dynamic memory inference networks provide a mechanism for sustained context recovery and updation transferable beyond language applications to sequential retail events [36]. Model governance will require pre-registered lists of features, verifications of leakages, and interventional validation on holdouts; interpretability will utilize Shapley-based summaries constrained to monotone business rules. Deployment will consist of offline policy evaluation through inverse-propensity weighing and online probing with constrained contextual bandits, setting safety ceilings on return rates and delivery latencies.

Policy Experiments

Pricing and zoning policies both set emissions and cost-to-serve. Subsequent versions will add randomized, verifiable experiments at the orchestration layer: real-time slot fees based on marginal emissions and route density; geofenced service territories that are toggled on or off with low-emissions-zone hours and congestion; and variable return windows tied into the item resale curves [33]. Randomization will occur geography-by-day or route-by-shift, with checks for covariate balance and CUPED variance reduction.

Each policy cell will capture intent-to-treat, compliance, and outcomes at the order, route, and facility level so that both spillover and direct impacts can be estimated. Monitoring over time will utilize alpha-spending to avoid peeking bias; the minimum detectable impact will be pre-calculated for the primary outcomes (margin per order and kg CO2e per order). Fail-safe mechanisms will automatically turn back policies when leading indicators cross guardrails, and an audit trail will link decisions and data snapshots along with model versions for internal and external assurance purposes.

CONCLUSION

This study shows that a joint profit–carbon lens is decision-useful and practical when grounded on order-level data and valid identification. With the combination of activity-based cost-to-serve with lifecycle carbon accounting and channel effect estimation via staggered difference-in-differences, event-study dynamics, and instrumental variables, channel design and secular demand are isolated in explaining contribution margin, ROIC, and order-weighted kilograms of CO_2-equivalent. Breakdown through transport, facilities, packaging, and return identifies what drivers induce observed gaps. Break- even surfaces identify where delivery, store trips, or BOPIS dominate profit and carbon. Heterogeneity through density, basket size, vehicle mode, and return propensity accounts for cross-market variation.

Insights translate into a playbook that unites telemetry and finance with operations. Route density, time-windows, and inventory positioning become first-order cost and footprint levers. Electric van and micromobility do best where stop density is great and curb accessibility is good; ship-from-store is preferred where short stems and breadth on-hand are dominant, and ship-from-distribution centers dominate where long-haul consolidation is prominent. Embedding a shadow carbon price into slot fees and choreography aligns customer decisions with minimal-emission tours. Carbon-adjusted contribution margin produces a governance metric harmonizing financial goals and lifecycle impact and guiding priorities.

The article advances practice through a reusable data architecture that harmonizes ledgers, OMS/WMS/TMS events, facility meters, and telematics into slowly changing dimensional canonical facts. State-of-the-art staggered-treatment estimators and panel dynamics respect treatment timing and persistence; instrumental-variables designs address simultaneity where the channel intensity is determined endogenously. Robustness is systematic: pre-trend checks, placebo groups, leave-one-firm-out diagnostics, tail return decile trimming, randomization inference, and cluster-robust uncertainty boost credibility. Structural decomposition via Oaxaca–Blinder and Shapley methods fills gaps for transport, facility energy, packaging, and return actions. Break-even charts and waterfalls convert coefficients into levers that managers can price, schedule, and monitor. Manage and measure at order granularity, and intervene first where the mediated effect is the most significant—engineer density through synchronized cutoffs, delivery-day consolidation, pickup, and locker locations. Co-optimize the placement of inventory and the last-mile routing such that availability and marginal delivery cost meet. Tune slot fees to reflect the incremental cost and marginal kilograms of CO_2-Eq to encourage customers to adopt the most efficient schedule. Standardize packaging with fragility-sensitive, cube-minimizing rules; charge damage-rate penalties for preventing the wrong kind of savings. Redesign the returns process to deter unnecessary reverse miles without compromising customer trust. Govern through dashboards that reveal stops per route-hour, grams of CO_2e per stop, and carbon-adjusted contribution margin per market.

Interpretation is limited by data quality and context. Third-party carrier reports can obscure load factors and lanes; facility energy can be coarsely metered; and packaging bills of materials are unavailable for legacy SKUs. These missing values require transparent imputation, propagated uncertainty, and prudent external disclosure. Generalizability relies on urban form, curb access, and grid composition; results for dense, temperate markets must be tested and established for sparse or hot markets with extended stems and cooling loads. Endogeneity from residuals is adjusted but not removed; sensitivity to definitions, thresholds, and estimator choices must be routinely examined and reported. The way forward is measurable and executable. Augmenting telemetry enables marginal-intensity routing and real-time facility scheduling; multi-echelon optimization with an internal carbon price aligns inventory with low-carbon fulfillment; and causal machine learning focuses the offers, the slots, the packaging, and the returns where the payoff is highest. Embedding auditable experiments—dynamic slot pricing, zonal availability, variable return windows—keep the changes statistically justifiable and reversible with safety. Managers can set guardrails on the grams CO_2e per stop and the carbon-adjusted contribution margin by market and iterate with order-level ledgers for money and for carbon. When executed well, firms push the profit–planet frontier outward, replacing the trade-offs with compounding, system-wide improvements.

REFERENCES

1. Abebe, F. “Factors Affecting the Modal Choice in Urban Mobility in Addis Ababa: The Case of Yeka Sub City.” 2022.
2. Baschetti, F., G. Bormetti, and P. Rossi. “Deep Calibration with Random Grids.” *Quantitative Finance*, vol. 24, no. 9, 2024, pp. 1263-1285.
3. Bassi, V., M. E. Kahn, N. L. Gracia, T. Porzio, and

J. Sorin. *Jobs in the Smog: Firm Location and Workers' Exposure to Pollution in African Cities*. National Bureau of Economic Research, 2022, w30536.

4. Boďa, M. *Methodology of Industry Statistics: Averages, Quantiles, and Responses to Atypical Values*. 2020.
5. Brissman, A., and T. Watson. "Auditors' Risk- and Materiality Assessment During Times of High Exogenous Risk: A Quantitative Study Analysing the Effects of the COVID-19 Pandemic." 2023.
6. Chavan, A. "Fault-Tolerant Event-Driven Systems: Techniques and Best Practices." *Journal of Engineering and Applied Sciences Technology*, vol. 6, 2024, E167. http://doi.org/10.47363/JEAST/2024(6)E167
7. Chavan, A., and Y. Romanov. "Managing Scalability and Cost in Microservices Architecture: Balancing Infinite Scalability with Financial Constraints." *Journal of Artificial Intelligence & Cloud Computing*, vol. 5, 2023, E102. https://doi.org/10.47363/JMHC/2023(5)E102
8. Chiu, I. H. "The EU Sustainable Finance Agenda: Developing Governance for Double Materiality in Sustainability Metrics." *European Business Organization Law Review*, vol. 23, no. 1, 2022, pp. 87-123.
9. Dehghanimadvar, M. *Multi-objective Optimization of PV Module Assembly Supply Chain, Concerning Economic, Environmental and Social Impacts*. PQDT-Global, 2023.
10. Dhanagari, M. R. "MongoDB and Data Consistency: Bridging the Gap between Performance and Reliability." *Journal of Computer Science and Technology Studies*, vol. 6, no. 2, 2024, pp. 183-198. https://doi.org/10.32996/jcsts.2024.6.2.21
11. Dhanagari, M. R. "Scaling with MongoDB: Solutions for Handling Big Data in Real-Time." *Journal of Computer Science and Technology Studies*, vol. 6, no. 5, 2024, pp. 246-264. https://doi.org/10.32996/jcsts.2024.6.5.20
12. Frank, K. A., et al. "Quantifying Sensitivity to Selection on Unobservables: Refining Oster's Coefficient of Proportionality."
13. Gorst, J., E. Ballantyne, and N. Dodd. "Measuring CO_2 Emissions from Transportation in a UK Retail Returns Supply Chain." *Journal of Supply Chain Management, Logistics and Procurement*, vol. 7, no. 2, 2024, pp. 194-206.
14. Jokinen, N. "Invoicing of Storage Costs for Ready-to-Deliver Goods in Case of Delayed Pick-Up from the Customer: Encouraging Collection of Finished Goods." 2024.
15. Kaakkunen, N. "Developing Customer Profitability Analysis from the Perspective of Fixed Costs and Depreciation Allocation." 2023.
16. Karwa, K. "AI-Powered Career Coaching: Evaluating Feedback Tools for Design Students." *Indian Journal of Economics & Business*, 2023. https://www.ashwinanokha.com/ijeb-v22-4-2023.php
17. Karwa, K. "Navigating the Job Market: Tailored Career Advice for Design Students." *International Journal of Emerging Business*, vol. 23, no. 2, 2024. https://www.ashwinanokha.com/ijeb-v23-2-2024.php
18. Kemp, N. J., et al. "Carbon Footprint of Alternative Grocery Shopping and Transportation Options from Retail Distribution Centers to Customer." *Environmental Science & Technology*, vol. 56, no. 16, 2022, pp. 11798-11806.
19. Khan, M. A., S. West, and T. Wuest. "Midlife Upgrade of Capital Equipment: A Servitization-Enabled, Value-Adding Alternative to Traditional Equipment Replacement Strategies." *CIRP Journal of Manufacturing Science and Technology*, vol. 29, 2020, pp. 232-244.
20. Konneru, N. M. K. "Integrating Security into CI/CD Pipelines: A DevSecOps Approach with SAST, DAST, and SCA Tools." *International Journal of Science and Research Archive*, 2021. Retrieved from https://ijsra.net/content/role-notification-scheduling-improving-patient
21. Kouadou, E. G. A. "The Impact of Financial Indicators on the Valuation and Investment Strategies of Industrial Companies: SOCOMEC'S Case." 2023.
22. Kripfganz, S., and D. C. Schneider. "ardl: Estimating Autoregressive Distributed Lag and Equilibrium Correction Models." *The Stata Journal*, vol. 23, no. 4, 2023, pp. 983-1019.
23. Kumar, A. "The Convergence of Predictive Analytics in Driving Business Intelligence and Enhancing DevOps Efficiency." *International Journal of Computational Engineering and Management*, vol. 6, no. 6, 2019, pp. 118-142. Retrieved from https://ijcem.in/wp-content/uploads/THE-CONVERGENCE-OF-PREDICTIVE-ANALYTICS-IN-DRIVING-BUSINESS-INTELLIGENCE-AND-ENHANCING-DEVOPS-EFFICIENCY.pdf
24. Kumar, S. M., and R. Venkatachalam. "Poor Economics and Its Missing Mechanisms: The Case for Causal Mediation." *Review of Development Economics*, vol. 28, no. 4, 2024, pp. 2014-2033.
25. Kusz, D., et al. "Determinants of Return-on-Equity (ROE) of Biogas Plants Operating in Poland." *Energies*, vol. 16, no. 1, 2022, p. 31.
26. Lee, K., J. Ko, and S. Jung. "Quantifying Uncertainty in Carbon Emission Estimation: Metrics and Methodologies." *Journal of Cleaner Production*, vol. 452, 2024, p. 142141.
27. Leung, M. P. "Cluster-Randomized Trials with Cross-Cluster Interference." arXiv preprint arXiv:2310.18836, 2023.
28. Malik, G. "Business Continuity & Incident Response." *Journal of Information Systems Engineering and Management*, vol. 10, no. 45s, 2025, pp. 451–473. https://www.jisem-journal.com/index.php/journal/article/view/8891
29. Malik, G., and Prashasti. "Implementing Zero Trust Architecture: Modern Approaches to Secure Enterprise Networks." *International Journal of Networks and Security*, vol. 5, no. 1, 2025, pp. 22–

45. https://www.academicpublishers.org/journals/index.php/ijns/article/view/3734

30. Mawson, V. J., and B. R. Hughes. "Optimisation of HVAC Control and Manufacturing Schedules for the Reduction of Peak Energy Demand in the Manufacturing Sector." *Energy*, vol. 227, 2021, p. 120436.
31. McBride, J. G. *Race, Suburbs, and the Municipal Bond Market: Evidence from California.* University of California, Los Angeles, 2024.
32. Menini, C., and S. Ruozzo. "Outsourcing Strategic Choices in the Middle Mile Network for a Merchant in B2C E-Commerce: A Comprehensive Decision Framework." 2021.
33. Munir, T. *Development and Evaluation of Road Network Pricing Models for Travel Demand Management*. PhD dissertation, Swinburne University, 2024.
34. Natarajan, T., and D. R. Veera Raghavan. "Does Pickup Service Quality Explain BOPIS Users' Store Relationship Performance? The Moderating Role of Users' Service Experience Consciousness." *The TQM Journal*, vol. 36, no. 8, 2024, pp. 2418-2448.
35. Papaioannou, V. *Uncertainty in GB Electricity Grid Carbon Intensity and Its Implications for Carbon Accounting and Reporting*. PhD dissertation, University of Reading, 2020.
36. Raju, R. K. "Dynamic Memory Inference Network for Natural Language Inference." *International Journal of Science and Research (IJSR)*, vol. 6, no. 2, 2017. https://www.ijsr.net/archive/v6i2/SR24926091431.pdf
37. Rodríguez-Garcia, M., A. O. Bas, J. C. Prado-Prado, and A. Lyons. "Fulfillment Costs in Online Grocery Retailing: Comparing Retail Store and Warehouse Strategies." *International Journal of Production Management and Engineering*, vol. 11, no. 2, 2023, pp. 127-145.
38. Rossolov, A., H. Rossolova, and J. Holguín-Veras. "Online and In-Store Purchase Behavior: Shopping Channel Choice in a Developing Economy." *Transportation*, vol. 48, no. 6, 2021, pp. 3143-3179.
39. Sardana, J. "Scalable Systems for Healthcare Communication: A Design Perspective." *International Journal of Science and Research Archive*, 2022. https://doi.org/10.30574/ijsra.2022.7.2.0253
40. Schnieder, M., C. Hinde, and A. West. "Sensitivity Analysis of Emission Models of Parcel Lockers vs. Home Delivery Based on HBEFA." *International Journal of Environmental Research and Public Health*, vol. 18, no. 12, 2021, p. 6325.
41. Silva, V., A. Amaral, and T. Fontes. "Sustainable Urban Last-Mile Logistics: A Systematic Literature Review." *Sustainability*, vol. 15, no. 3, 2023, p. 2285.
42. Singh, V. "Advanced Generative Models for 3D Multi-Object Scene Generation: Exploring the Use of Cutting-Edge Generative Models like Diffusion Models to Synthesize Complex 3D Environments." 2022. https://doi.org/10.47363/JAICC/2022(1)E224
43. Singh, V. "Explainable AI in Healthcare Diagnostics: Making AI Models More Transparent to Gain Trust in Medical Decision-Making Processes." *International Journal of Research in Information Technology and Computing*, vol. 4, no. 2, 2022. https://romanpub.com/ijaetv4-2-2022.php
44. Torchio, F. *Survey on Automated Systems for Smart Warehouses*. PhD dissertation, Politecnico di Torino, 2023.
45. Umanetz, A. *Predicting Fleet-Vehicle Energy Consumption with Trip Segmentation*. PhD dissertation, 2021.
46. Wang, S. *Novel Optimization-based Methods for Downscaling Power Sector Emission Projections and Compensating Reliability Contributions of Power Supply*. PhD dissertation, Johns Hopkins University, 2023.
47. Yang, Y., and S. Liu. "Asset Shortages and Economic Operating System Based on Two-Stage Least-Squares Regression." *International Journal of Shipping and Transport Logistics*, vol. 13, no. 5, 2021, pp. 466-483.