DESIGNING A MARKET DECISION STACK FOR FAIR BUDGET ALLOCATION IN INDIA FMCG

A Thesis Submitted
In Partial Fulfillment of the Requirements
for the Degree of

EXECUTIVE MASTER OF BUSINESS ADMINISTRATION

in

Data Science Analytics

by

SATYENDRA
(Enrollment No. 24/UEMBA/09)

Under the Supervision of

Dr. Deepti Aggrawal
University School of Management and Entrepreneurship (USME)
Delhi Technological University

To the University School of Management and Entrepreneurship (USME)

DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-110042, India

July, 2026

ACKNOWLEDGEMENTS

I express my sincere gratitude to my supervisor, **Dr. Deepti Aggrawal**, for guidance throughout this thesis work.

I thank the faculty and staff of the University School of Management and Entrepreneurship (USME), Delhi Technological University, for academic support.

I acknowledge the public data agencies whose published statistics underpin this secondary-data Decision Stack (MoSPI, RBI, TRAI, Census of India / NFHS, and related portals cited in Appendix-I and the References).

Any remaining errors are my own.

Satyendra

Satyendra

CANDIDATE'S DECLARATION

I **Satyendra** hereby certify that the work which is being presented in the thesis entitled **Designing a Market Decision Stack for Fair Budget Allocation in India FMCG: State–category priority and fair shares from public secondary data** in partial fulfillment of the requirements for the award of the Degree of **Executive Master of Business Administration in Data Science Analytics**, submitted in the University School of Management and Entrepreneurship (USME), Delhi Technological University is an authentic record of my own work carried out during the period from **June 2026** to **July 2026** under the supervision of **Dr. Deepti Aggrawal**.

The matter presented in the thesis has not been submitted by me for the award of any other degree of this or any other Institute.

Satyendra

Candidate's Signature

This is to certify that the student has incorporated all the corrections suggested by the examiners in the thesis and the statement made by the candidate is correct to the best of our knowledge.

Signature of Supervisor(s)

Signature of External Examiner

CERTIFICATE

Certified that **Satyendra** (enrollment no. **24/UEMBA/09**) has carried out their research work presented in this thesis entitled "Designing a Market Decision Stack for Fair Budget Allocation in India FMCG: State–category priority and fair shares from public secondary data" for the award of **Executive Master of Business Administration in Data Science Analytics** from the University School of Management and Entrepreneurship (USME), Delhi Technological University, Delhi, under my supervision. The thesis embodies results of original work, and studies are carried out by the student himself and the contents of the thesis do not form the basis for the award of any other degree to the candidate or to anybody else from this or any other University/Institution.

	Signature **(Dr. Deepti Aggrawal)** Supervisor University School of Management and Entrepreneurship (USME), Delhi Technological University

Date: 30.07.2026.

ABSTRACT

The FMCG national budget for India must be allocated to state-category cells which are not homogeneous. Affordability, digital preparedness, CPI pressure, festival dates, and advertisement clutter differ. Suitability charts provide us with regions that present opportunity and money flows there. In this study, we end this process with a constrained ₹ allocation an examiner can audit, and with a minimum allocation threshold that will not exclude smaller markets from consideration.

Public data sources come in as a harmonized secondary panel by state × category × month. Opportunity Index (OI) measures structural suitability. Growth propensity captures the short-term dynamics. Finally, logistic regression, Random Forest, XGBoost, and SVM algorithms compete in a time cross-validation. ROC-AUC and average precision select a winner among them. The combined scores determine Investment/Protect/Probe/Suppress decision. An illustrative budget is sliced into EV share values. Allocation fairness thresholds and digital caps follow. One experiment writes the playbook and the thesis. Panel data goes to Data & Appendix-I.

The linear model defeated the tree learners in the bake-off for holdout. In instances where OI and propensity were congruent, the Invest cells consumed a big chunk of the sample budget. The floors pushed expenditure towards more restricted eligible markets. The EV ranking order remained the same. The counts and rupee differences are in Results.

The board receives an audit trail from total public numbers to priority positions and restricted Rupee allocations, cell by cell. Up next is to anchor portal files to actual releases and sharper category testing.

Keywords: ** India FMCG, geo-marketing, Opportunity Index, growth propensity, logistic regression, model comparison, fairness-constrained budget allocation, secondary data, decision analytics.

LIST OF TABLES

Table	Caption (content)
Table 1.1	Unit of analysis used in this study
Table 1.2	Research objectives and corresponding Decision Stack stages
Table 1.3	Design hypotheses H1–H12 — claimed direction and Decision Stack implementation
Table 2.1	PRISMA screen counts by stage (identification through narrative-synthesis core)
Table 2.2	Theory-strand coverage in the Scopus synthesis core versus Scholar/OA gap-fill
Table 2.3	Demand constructs and public-panel proxies used in this study
Table 2.4	TCCM synthesis — dominant findings and implications for this study
Table 2.5	ADO synthesis — antecedents, decisions, and outcomes mapped to the Decision Stack
Table 2.6	What each theory strand establishes, and what remains missing for India FMCG planners who must work from public data
Table 3.1	Design rules that govern every data join in this study
Table 3.2	Researcher terminology for the secondary panel used in this report
Table 3.3	Data layers, public sources, panel roles, and reference pointers
Table 3.4	Included constructs versus commonly suggested alternatives that were set aside
Table 3.5	Illustrative CPI / public-group map to six FMCG decision categories
Table 3.6	Processed panel schema — core columns and meanings
Table 3.7	Panel coverage metrics from the reported pipeline run
Table 3.8	Known frictions when raw portal downloads replace harmonised joins
Table 4.1	Propensity-model feature set and roles
Table 4.2	Pre-specified CPI lag features — construction and role
Table 4.3	Priority cells from median splits of Opportunity Index × growth propensity
Table 4.4	Sensitivity scenarios — construction and managerial reading
Table 4.5	Method choices retained for this study, with the main alternative set aside
Table 5.1	What the Decision Stack delivers for planning
Table 6.1	Pipeline scale — panel size, budget, assumed lift, and holdout months
Table 6.2	Propensity model bake-off — holdout ROC-AUC, average precision, and selection
Table 6.3	Selected propensity model metrics (logistic regression)
Table 6.4	Largest standardised logistic coefficients (absolute size)

Table	Caption (content)
Table 6.5	Continuous Ridge baseline versus logistic propensity on the same features and 3-month holdout
Table 6.6	Holdout-length bake-off and walk-forward stability (winner by ROC-AUC)
Table 6.7	Latest-month priority matrix — cell counts, mean OI, mean propensity, mean population
Table 6.8	Fairness-constrained recommended budget by priority cell (₹ crore and share of B)
Table 6.9	Fairness constraint metadata — eligible cells, floors, cutoffs, and reallocation
Table 6.10	Discriminative sensitivity — Invest budget and fairness-pool delta by scenario
Table 6.11	Design hypotheses H1–H12 — how this study treats each claim
Table 6.12	Results evidence map
Table 6.13	Panel reliability (missingness & cardinality) — live analytics
Table 6.14	OI × propensity priority matrix (latest-month cells) — live analytics
Table 6.15	Budget by PriorityCell — live analytics
Table 6.16	Fairness-constrained EV budget — top 40 cells — live analytics
Table 6.17	Lowest 15 budget cells (Suppress / Probe tail) — live analytics
Table 6.18	Sensitivity summary — live analytics
Table 6.19	Fairness constraint metadata — live analytics
Table 6.20	Selected model and run metadata (from model_metrics.json) — live analytics
Table 7.1	Managerial stance by priority cell
Table 8.1	Key risks and mitigations applied in this study
Table 8.2	Social impact pathways linked to Decision Stack outputs on the reported run
Table I.1	Dataset designation terms, definitions, and status in this bind
Table I.2	Analytical uses, variables, publishers, cites, URLs, and retrieval dates
Table I.3	Coverage dimensions of the harmonized secondary panel
Table I.4	Panel-layer inventory — encodings and source basis
Table I.5	CPI lag design features — construction and justification
Table I.6	Nested holdout lag-spec comparison (ROC-AUC and average precision)
Table I.7	Defence diagnostics summary (latest-month surface)
Table I.8	Tier-2 cross-check summary (matches Tables 6.5–6.6)

CONTENTS

Title	Page No.
Abstract	v
List of Tables	vi
List of Figures	ix
List of Symbols, Abbreviations and Nomenclature	x
CHAPTER 1: INTRODUCTION	1
CHAPTER 2: LITERATURE REVIEW	5
CHAPTER 3: DATA	16
CHAPTER 4: METHODOLOGY	21
CHAPTER 5: NOVELTY AND CONTRIBUTION	27
CHAPTER 6: RESULTS AND FINDINGS	29
CHAPTER 7: MANAGERIAL IMPLICATIONS	66
CHAPTER 8: CONCLUSION, FUTURE SCOPE AND SOCIAL IMPACT	68
REFERENCES	71
APPENDIX-I: DATA SOURCES, RETRIEVAL, AND PANEL CONSTRUCTION	74
Plagiarism Verification	81

LIST OF FIGURES

Figure	Caption
Fig. 2.1	PRISMA literature screen flow
Fig. 2.2	Data variables → predictive dependent variable → decision dependent variables
Fig. 6.1	Propensity model bake-off (logistic / RF / XGBoost / SVM)
Fig. 6.2	Market Decision Stack analytic flow
Fig. 6.3	OI × growth propensity by priority tier
Fig. 6.4	OI × propensity priority matrix (latest month)
Fig. 6.5	Fairness-constrained budget by priority tier
Fig. 6.6	Largest fairness lifts and cuts
Fig. 6.7	Top recommended state × category cells
Fig. 6.8	OI distribution by priority tier
Fig. 6.9	Category × priority cell counts
Fig. 6.10	Budget before vs after fairness
Fig. I.1	CPI growth ACF/PACF (lag diagnostics)
Fig. I.2	Nested lag-spec holdout AUC

LIST OF SYMBOLS, ABBREVIATIONS AND NOMENCLATURE

Symbol / Abbreviation	Meaning
AP	Average precision
AUC / ROC-AUC	Area under the receiver operating characteristic curve
	Illustrative national marketing budget (₹ crore)
CPI	Consumer Price Index
EV	Expected value (budget share)
FMCG	Fast-Moving Consumer Goods
GSDP	Gross State Domestic Product
	Assumed response lift (planning assumption, not estimated elasticity)
MoSPI	Ministry of Statistics and Programme Implementation
NFHS	National Family Health Survey
OI	Opportunity Index
RBI	Reserve Bank of India
RF	Random Forest
TRAI	Telecom Regulatory Authority of India
UIDAI	Unique Identification Authority of India

CHAPTER 1
INTRODUCTION

1.1 Context

The FMCG surface of India is uneven across geography. While food essentials may continue to seem sparse in Bihar, it is highly saturated and technologically sophisticated in Maharashtra. A single national media spend item, but two very different numbers.

Groups like CPI, GSDP per capita, TRAI-like wireless and broadband preparedness, festivals, and advertising activity all come together into such a surface. Geo-marketing already provides location scores and suitability. However, planners still require a mapping principle that maps do not give – how those scores become allocations of a national total budget (B) where volume constraints and floors apply alongside. Inclusive growth and Bharat-first considerations frequently impose floors on thin against thick geographies. Growth tendency and EVs share this connection with such floors on one surface.

1.2 Problem and business question

A national budget constraint has many combinations to consider at once between state and category. Cost, digitization, and ad pollution are important. There should also be a fairness floor in the thin markets case. What is not included? The spend rules themselves. What cells should receive spend, and in what fraction ₹?

Business question: for (B) = ₹10 crore national budget, what state-category cells get spend, and how much?

The approach chosen for this task is Market Decision Stack on public panel data state x category x month. Dataset information can be found in Chapter 3 and Appendix-I.

1.3 Unit of analysis

Markets are scored at state × category × month. Table 1.1 lists the four dimensions and what each does inside the Decision Stack.

Table 1.1. Unit of analysis used in this study.

Dimension	What is measured	Role in the Decision Stack
State	Indian state / UT on a canonical code table	Geography for national budget splits. Matches public MoSPI / RBI / TRAI-style releases
FMCG decision category	One of six planning categories	Separates staples, beverages, personal care, household care, packaged foods, and oral care inside the same state
Month	Calendar month in a multi-year panel	Captures CPI movement, festivals, and readiness over time. Supports time-holdout propensity and latest-month ₹ shares
Data basis	Public secondary aggregates	Joined to the state × category × month schema (Chapter 3)

The national budget itself is already divided according to the Indian geography. So are the MoSPI, RBI, and TRAI releases. State, then, is the natural grain. Within a state, categories differ. Food_Staples may be healthy when Personal_Care is saturated. There are six planning categories that come next: Food_Staples, Beverages, Personal_Care, Household_Care, Packaged_Foods, and

Oral_Care. Inflation, festival season, and preparedness change on a monthly basis. Every line represents a state × category × month data point. The allocator locks in the most recent month when it sets the ₹ share.

1.4 Research objectives

The objectives follow the Scopus geomarketing review (PRISMA core), Scholar/OA gap notes, and the packaging gap from the lit review. Table 1.2 states the four research objectives and the Decision Stack stage each one advances.

Table 1.2. Research objectives and corresponding Decision Stack stages.

Objective	What this study does	Decision Stack stage
O1	Construct an India FMCG Decision Stack on free public secondary panels at state × category × month (sources and retrieval in Chapter 3 / Appendix-I)	Panel construction
O2	Score structural Opportunity Index (OI) from unmet demand, affordability, channel readiness, and ad clutter. Estimate near-term growth propensity through a classifier bake-off (logistic regression, Random Forest, XGBoost, SVM) on a time-based holdout selected by ROC-AUC and average precision. Rank cells on an OI × propensity matrix (Invest / Protect / Probe / Suppress)	Opportunity, propensity, and priority
O3	Convert priority cells into illustrative expected-value ₹ shares, then apply India-specific fairness floors and digital-share caps, with sensitivity tables	Allocation and constraints
O4	Use festival seasonality as propensity features and as an allocation sensitivity check	Seasonality check

Table 1.2 therefore runs from panel construction through scoring and priority to constrained ₹ shares. Festival effects stay visible twice. They enter once as propensity features and once as a stress test.

1.5 Research hypotheses

The three strands of literature that make up the stack are geosuitability, marketing measures/budget allocation, and inclusive growth practice. This is not something that has been reproduced using H1-H12. Rather, what is provided are concrete design actions. The structural opportunity increases due to residual demand, ability to pay, and channel capacity, but decreases on account of clutter. Growth propensity for the short term is determined by the demand persistence, festival pressure, and digital readiness. This is kept independent of the structural score. The budget allocation follows accordingly.

Twelve design hypotheses H1 to H12 operationalize those moves into propositions that can be audited by an examiner on this panel. These hypotheses derive from the synthesis presented in Chapter 2 as Decision Stack principles or directional predictions. These are not independent theories in and of themselves, nor do they represent paths estimated by SEM. Hold-out ROC-AUC and average precision confirm the propensity ranking for growth momentum.

Table 1.3. Design hypotheses H1–H12 — claimed direction and Decision Stack implementation.

Block	Hypothesis	Claimed direction	Implemented as
Opportunity formation	H1	Higher unmet demand raises Spatial Opportunity (Opportunity Index)	Equal-weight within-month OI term
Opportunity formation	H2	Higher affordability raises Spatial Opportunity	Equal-weight within-month OI term
Opportunity formation	H3	Higher digital and channel readiness raises Spatial Opportunity	Equal-weight within-month OI term
Opportunity formation	H4	Higher advertising clutter lowers Spatial Opportunity	Within-month OI clutter penalty
Growth propensity	H5	Festival pressure raises growth propensity	Propensity feature
Growth propensity	H6	Demand persistence (lagged CPI and year-on-year movement) raises growth propensity	Propensity feature
Growth propensity	H7	Digital readiness raises growth propensity	Propensity feature
Growth propensity	H8	Spatial Opportunity and Growth Propensity remain distinct constructs on a two-axis decision surface	Separate axes until priority
Priority and allocation	H9	Higher Opportunity Index and higher growth propensity together raise the chance of an Invest posture	Median-split priority matrix
Priority and allocation	H10	Expected-value response guides unconstrained budget share	EV allocator
Priority and allocation	H11	Fairness pressure raises floor-adjusted share for lagging (low-penetration / low-GSDP) cells	Fairness floors
Priority and allocation	H12	Low digital readiness lowers the allowable digital share of budget through a hard cap	Digital-share cap

H1–H4 are realised in the Opportunity Index. H5–H7 sit in the propensity feature set. H8–H9 use the two-axis priority matrix. H10–H12 use the expected-value allocator with fairness floors and digital caps. Chapter 6 maps each claim to empirical evidence.

1.6 Scope

The working model is a Decision Stack for India FMCG markets using publicly available data. Covered in scope are panel construction (Chapter 3), the Opportunity Index, time-holdout growth propensity, Invest/Protect/Probe/Suppress cells, expected-value allocation with fairness floors and digital share ceilings, and the twelve design hypotheses in Table 1.3. The fairness floors remain an adjustable corporate policy decision, and are not assumed to be socially optimal.

What is new is the packaging. Known approaches are baked into open priority and budgeting rules. Hypotheses H1-H12 are design hypotheses drawn from the literature review. Growth propensity is tested primarily using holdout AUC/AP, and not by SEM path models. Outside the scope lies causal

media ROI, byte-identical dumping of portals for all series, shopping microdata, scanner data, and brand creative testing.

Literature and gap come next (Chapter 2). Then the working panel (3), methods (4), novelty honesty (5), results (6), and managerial guidance (7). Limitations, future scope, and social impact close the bind in Chapter 8.

CHAPTER 2
LITERATURE REVIEW

2.1 Purpose of the review

Allocation of FMCG spending budget for the whole nation into Indian states and categories involves multiple streams of literature together. Geo-marketing uses place. Marketing metrics favor response-weighted spending. Ability to pay is influenced by demand and affordability. Digital capability affects the capability of channels to deliver the budget allocation. Practice of inclusive growth supports floor values for thin geographies. Classifiers and z-score opportunity indices have been established as methods. Response weighted budgeting has been established as well. It is the connection that this chapter needs to defend.

The review moves in four uneven steps so a single template does not hide where the strands diverge.

1. A transparent, PRISMA-adapted screen of the geomarketing literature (search scope and retention).
2. A thematic narrative synthesis around the stack's building blocks. It covers where opportunity sits, how near-term momentum is read, how budgets track expected response, and how fairness and seasonality enter Indian planning.
3. A short cross-cutting TCCM / ADO coding of the synthesis core, then an explicit research-gap statement that motivates the study objectives.
4. A conceptual framework that puts those strands on one auditable Market Decision Stack and ends with the data-to-decision map in Figure 2.2.

2.2 Search method

PRISMA 2020–based screening has been done for two Scopus exports (extracted on 2026-07-21). There was a broad geo/allocation sweep (43 records) and a geomarketing sweep (202 records). After deleting 1 duplicate record, 244 records have been screened by title and abstract and assessed for their adherence to the inclusion criteria. The inclusion criteria include a marketing/retailing decision-making environment and at least one Decision Stack strand.

67 records have been rejected as off-topic, namely cloud/satellite studies, natural resources GIS studies, conference reviews, and similar noise. 177 records proceeded to the eligibility stage. 10 records have failed the Decision Stack fit gate. 167 records remained eligible.

Out of the list of available articles, 15 Scopus articles are chosen for narrative synthesis. 7 important handpicked articles are included along with them (Bronnenberg & Mahajan, Farris et al., Lilien et al., Rossi et al., Tellis, Kotler et al. South Asia, NITI Aayog 2018).

Scholar / OpenAlex OA gap fill added 6 more texts on 2026-07-26. They helped fill the gaps in the theoretical strands present in the Scopus geomarketing core. Those strands were spatial models, budgeting of expected response, ADP fairness, and India festival seasonality.

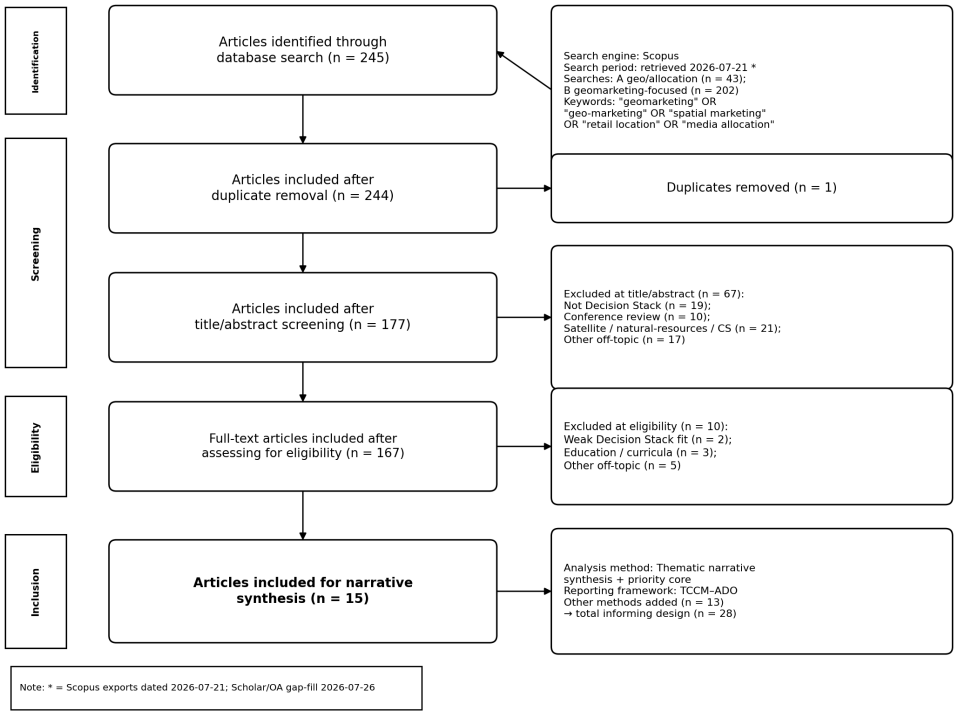
The 2026-07-31 methods OA fill included 6 additional anchors. These include equal weighting, sensitivity to composite indicators (OECD/JRC, Poledníková, Opportunity Index Technical Supplement, Zhao et al. GIS-MCE). Constrained advertising fairness (Celis et al., Gelauff/Ali et al.) was also included among them. The design corpus is now at 34 references (15 + 7 + 6 + 6).

Figure 2.1 and Table 2.1 keep the Scopus-screen counts. The methods fill sits under "other methods." It does not re-open identification. The thematic synthesis follows in Sections 2.3–2.8.

The protocol is adapted for a narrative review. It justifies study design (Page et al., 2021, PRISMA 2020 logic).

2.2.1 PRISMA flow diagram

PRISMA 2020 adapted flow diagram for this study's literature screen



PRISMA 2020 adapted flow diagram for this study's literature screen

Figure 2.1. PRISMA 2020–adapted flow. Scopus ID → screen → fit check → core set + other methods. Identified 245. Other methods at figure freeze 13 (7 seminal + 6 Scholar/OA gap-fill). Dupes removed 1. After dedupe 244. Title/abstract out 67. Kept 177. Fit check out 10. Fit 167. Scopus core 15. Table 2.1 adds the later methods OA fill. Total design corpus 34.

Table 2.1. PRISMA screen counts by stage (identification through narrative-synthesis core).

Stage	n
Identified (Scopus)	245
Other methods (seminal hand-selected)	7
Other methods (Scholar / OA gap-fill, 2026-07-26)	6
Other methods (methods OA fill, 2026-07-31)	6
Duplicates removed	1
After deduplication	244
Title/abstract excluded	67
Retained after title/abstract	177
Eligibility excluded	10
Eligible	167
Scopus narrative-synthesis core	15
Total informing Ch.2 design (15 + 7 + 6 + 6)	34

Strand coverage in the Scopus synthesis core points to the research-gap analysis later in this review. The gap-fill columns show where full-text anchors were added outside the Scopus geomarketing hit list:

Table 2.2. Theory-strand coverage in the Scopus synthesis core versus OA gap-fills.

Theory strand	In Scopus synthesis core	After Scholar/OA and methods fills
Geo-marketing / multimarket spatial allocation	15	+ Bradlow et al. (2005) spatial-models review
Marketing metrics / media–budget allocation	2	+ Gupta & Steenburgh (2008), textbooks remain
FMCG elasticity / affordability	2	Still thin — Bawindsom + Kotler South Asia
Digital readiness / channels	8	Covered (Garcia, Bawindsom, Crisóstomo, Kita)
Inclusive growth / fairness floors	0	+ NITI 2018 + NITI 2023 ADP, + Celis et al. (2019), Gelauff/Ali et al. (2020) constrained fair ads
Festival / seasonality	0	+ NIPFP Pandey (2022), Tamilselvi & Rajeswari, Samak et al. (2026)
Composite opportunity weighting / sensitivity	0 in geomarketing core as methods texts	+ OECD/JRC (2008), Poledníková (2017), Opportunity Index Tech. Supp. (2023), Zhao et al. (2024) EW baseline

There are two key screen results relevant for design. Scopus geomarketing provides a complete package right now. It includes spatial market prioritization, retail fitness and opportunities assessment. Fairness floors and festival seasonality are crucial for Indian nation-wide planning processes. They are almost absent from that famous list.

Gap-fills comment each missing piece separately. Methods fill provides equal weight and constrained fairness approaches to the Opportunity Index and allocator. Further chapters combine all in one package. Opportunity, propensity, EV ₹ shares, fairness and festival considerations are combined on India public secondary panels.

2.3 Spatial marketing and geomarketing

Multimarket studies in the classical period proved that the spatial relationship exists between market shares and promotion techniques. Marketers and competitors produce linked results in different markets (Bronnenberg & Mahajan, 2001). Bradlow et al. (2005) discuss spatial models in marketing. They prove that spatial dependence is a modeling issue with actual impact on market scores.

This idea has been further developed by more recent works on geomarketing and spatial marketing. These works include GIS-assisted retail location and trade area analyses, as well as the use of a wider “geomarketing mix,” which connects location with promotion as well as mobile location (Cliquet, 2020; Cliquet & Baray, 2020; Baray, 2025). The theory of retail location has a long history, going back to gravity theory and central place theory (Reilly, Huff, Christaller), then to multiplicative competitive interaction. Location-allocation models can now be found in geomarketing software packages (Cliquet, 2025). There is bibliometric evidence that GIS has been used for location analysis and spatial segmentation. Its adoption into marketing theory is still incomplete (Tudor, Girlovan & Botoroga, 2025).

In relation to FMCG planning in India, the stream under consideration allows three design considerations. The markets need to be looked at as spatial entities. Category relationships within these locations are important. This is the reason why the unit of analysis used will be state x category. Spatial decision support needs to be transparent enough to be audited. The working grains in this case are state x category x month.

2.4 Opportunity identification

The second group of research identifies where opportunities exist. They typically employ fit or white spot analysis. They seldom translate such map into budget allocation share.

In his study of German retail fit, Tudor (2025) develops a free geospatial workflow. It employs 1 km census grids and points of interest from OpenStreetMap. The workflow combines multi-criteria fit modelling with Moran's I / LISA and spatial weighted regression. The gap analysis reveals 40 cells that have high fit but do not have retail establishment. These cover both urban and rural regions. This is one example of explicit opportunity identification as the output. It corresponds to structural Opportunity Index prior to spending rules application.

Tudor (2025) also recomputes fit with alternative weights including equal weighting. The fit surfaces remain highly correlated. This implies that equal weighting composites might serve as a robust default approach in the case if stakeholder AHP is not available. Salgado et al. (2019) demonstrate how two-step WebGIS approach (k-means, Voronoi, weighted midpoint, kernel density) can help with selection of parent store and branches. It consumes less resources than desktop GIS.

Kerzina, Antineskul and Patrusheva (2023) research geomarketing guidelines for FMCG "shop-at-home" positioning in Perm. They consider catchment by demographics. Wichmann, Scholdra and Reinartz (2021) relate brick-and-mortar presence to inner-city attractiveness. They reveal how store location contributes to overall place ratings.

Gladilin and Levina (2020) contribute a demand momentum perspective. Bank transactions and geoclustering estimate the demand for a product category across urban districts. The researchers recommend opening / closing stores or changing assortment. Their process anticipates the split between place attractiveness and category momentum that will come in later literature.

As the issue of multi-criteria opportunity scores weighting, one should consider the methodological problem. Geomarketing articles often ignore it. As OECD/JRC Handbook on Constructing Composite Indicators states, most composite indicators are calculated using equal weighting. Each of them can be considered equally important, or there is not enough agreement for different weightings. In any case, weight selection must be stress-tested in uncertainty / sensitivity analysis (OECD, 2008).

Conclusion: Poledníková (2017) describes the Occam basis. In the absence of weight agreement, a simple equal weight approach is justified. Applied opportunity research adheres to this principle as well. The United States Opportunity Index takes an average of four dimensions with equal weights at the state and county level (Opportunity Index Technical Supplement, 2023).

Equal weights in the context of multi-criteria analysis in GIS represent a neutral baseline that allows for a comparison to AHP or machine learning weight generation algorithms (Zhao et al., 2024). In this study, the equal within month z-weights for unmet demand, affordability, channel readiness, and clutter penalty have been used. Weights and shocks sensitivity remain within the working stack. Welfare optimum is not calculated for the four constructs.

Together, these papers provide the 'where to prioritize' kit. Grids for fit, spot detection, WebGIS locating and geoclustering of banks all tell us something. Good places can be found even without brand scanning. All that remains unanswered in a national FMCG budgeting meeting is what comes after. The opportunity is defined. How much of a finite ₹ budget does each grid receive? There are many constraints: affordability, clutter, readiness and equity. Later chapters will make the transition from opportunity grid to weighted ₹ allocation grid. So does the model described in Section 2.11.

2.5 Marketing metrics and budgeting

The principle of good marketing metrics practice states that budget allocations need to be aligned with anticipated response rates, not necessarily historic spending figures (Farris et al., 2010). According to Gupta and Steenburgh (2008), the process needs two steps: estimating demand/response first and allocating resources after. Such a process supports EV allocation chosen by me.

Marketing-engineering texts describe geographical and resource-allocation problems as decision models, making assumptions clear and testable (Lilien, Rangaswamy & De Bruyn, 2013). Literature on advertising responses describes the whole picture: what ads do, how fast they become inefficient and their contextual influence (Tellis, 2004). Hierarchical Bayesian demand models form the background for market mix (Rossi, Allenby & McCulloch, 2005). Usually, however, they require more complete brand panels than allowed in this study due to public data limitations.

Media–budget allocation within the Scopus geomarketing core is rather sparse. There are only two studies that address the issue of the proper link between anticipated response and spending in geographic space. Location-based marketing is considered primarily a media mix tool which sends offers "at the most opportune time and place." However, implementation friction is rather high, since it includes such aspects as privacy issues, relevancy and notification overload (Garcia,2022).

Early geomarketing rules for mobile ad delivery stress targeted place-based promo. National expected-value budgeting is rarely the focus (De Chiara et al., 2011). The Gupta & Steenburgh working paper (OA gap-fill) fills that gap. It gives a full-text allocation anchor. Scopus geomarketing hits rarely offer one.

On the public panel for this study, an assumed response lift (= 0.04) feeds an explicit EV form (Eqn. 2.1):

$$EV_i = n_i \times p_i \times OI_i^+ \times \ell \qquad (2.1)$$

where (n_i) is a population weight, (p_i) is growth propensity, and (OI^{+}_i) is a positive transform of opportunity. The formula is clear enough for board stress tests. Baseline shares come with sensitivity tables (CPI-like OI shock, festival filter). Planners can see how the surface moves when () or the opportunity layer shifts.

2.6 Affordability and demand

FMCG consumption depends on income and price levels. Indian managerial literature discusses prioritization geographically as one of the basic marketing decisions in changing conditions locally (Kotler, Keller, Koshy & Jha, 2013). As for the Scopus geomarketing core, affordability is only referenced through analogies. For example, Bawindsom et al. (2024) employ living standards index. This is the measure of affordability of the household when evaluating the priority of FTTH/FWA broadband deployment in urban Burkina Faso. Potential and ability to pay precede the choice of technology portfolio.

Kita and Čvirik (2024) reveal that the level of digital engagement differs between the formats of stores within food retailing sector. This example also illustrates that the company's capabilities define the possible spatial strategy.

Without structural price elasticity from scanner data, the proxy of demand variables includes panels from India:

Table 2.3. Demand constructs and public-panel proxies used in this study.

Construct	Proxy in this study
Price / demand pressure	MoSPI-style CPI group indices by state × mapped category × month
Ability to pay	GSDP per capita (MoSPI / RBI Handbook style)
Market size	State population (Census / UIDAI / NFHS summaries)

These series enter the stack as decision proxies for opportunity and propensity. That reading is kept when results are read.

2.7 Digital readiness

Media and Mix in India have now become more focused on digital penetration and e-commerce delivery. The Scopus filter reveals a nearby example that forms the digital readiness layer. Garcia (2022) reports on location-based/geofenced media. Crisóstomo-Berrocal, Sierra-Liñan and Carbonell-Michael (2022) report on the digital geomarketing platform for micro and small companies. Their pre–post results include labour productivity +30.86%. Bawindsom et al. (2024) report ML-GIS-based broadband deployment ranking. Cliquet (2020) reports the move from geomarketing to location-based marketing on mobile devices.

TRAI Performance Indicator Reports publish statistics on wireless and broadband communications. These are publicly available, citable and state-sponsored. They can be used as digital readiness priors for the panel. The manager-level message is clear. Digital share should follow readiness. For the low broadband readiness cells, the allocator sets a maximum digital share (default 25%). Expected value shares will remain valid.

2.8 Fairness and festival seasonality

Spatial inclusion and geographically balanced development appear in the Scopus screen as policy topics. They do not appear as geomarketing findings. NITI Aayog (2018) addresses inclusive growth, regional equality and programs like the Aspirational Districts. They address dualism and lagging regions. NITI Aayog (2023) Assesment of Aspirational Districts Program measures progress on this uplift agenda. It lends support to the policy rationale to treat lagging markets as fairness eligible in budget rules.

Ad adjacent ad-system literature makes the same formal case. Fairness appears as a constraint on the allocation problem. Unconstrained optimization is not taken to solve the equity problem in its own right. Celis, Mehrotra and Vishnoi (2019) examine revenue optimal ad auctions with coverage constraints for fairness. Using a large scale search ad dataset, type-wise uniform coverage is achievable with little sacrifice in revenue.

Gelauff, Goel, Munagala and Yandamuri (2020) consider fair advertising for groups. Campaign budgets are divided to ensure that parity constraints hold within the constraints of platform spending rules. This body of literature does not provide the India FMCG geography. However, it provides support for treating fairness constraints as limited means of planning with an open efficiency trade-off.

In marketing terms, a fairness floor is a constrained optimisation (Eqn. 2.2):

$$\max \sum_i EV_i \quad \text{s.t.} \quad \sum_i b_i = B, \quad b_j \geq \underline{b}\, \forall j \in \mathcal{F} \qquad (2.2)$$

where () is the fairness-eligible set (low GSDP and/or low category penetration). The setup uses a ruled reallocation (floors + renormalise). A black-box optimiser would hide audit trails. Every floor and renormalisation step here stays open to audit. Floor levels and eligibility percentiles are policy parameters that boards can stress-test. That aligns with ADP-style lagging-unit targeting and with constrained fair-ad practice.

The festival calendar of Diwali, Holi, Onam, Eid, and other regional crop festivals lifts staple consumption in a regular pattern. These festival calendars also lift adjacent packaged food items, personal care products, and home cleaning products. This pattern is well-known to Indian brand managers. But no research in the Scopus geomarketing literature synthesis core codes the festival seasonality as a term in spatial allocation.

Domain and OA literature cover the theoretical gap. Pandey (2022, NIPFP) considers Diwali a mobile holiday and discusses its role in shifting seasonal consumption without establishing a long-term market leadership position. Tamilselvi & Rajeswari describe festive sale purchases in Chennai. Samak et al. (2026) investigate buyer behaviour and competition in Diwali sweets segments.

These anchors shape the use of festival flags and strength features in the propensity model. These anchors also drive an optional festival cell filter in sensitivity allocation. The peak season capacity then gets into the planning surface. The festivals were not already framed in geomarketing and this research does not claim otherwise.

2.9 Cross-cutting synthesis

Thematic reading resulted in coding of the narrative core. Gap-fill full texts are provided. Theory building matrices of two types of business review analysis are applied (Paul & Criado matrix / Paul & Rosado-Serrano matrix). TCCM (Theory–Context–Characteristics–Methodology) illustrates how the corpus is contextualized and supported. ADO (Antecedents–Decisions–Outcomes) reveals the manager sequence that a geo-budget stack needs to be made auditable by. Table 2.4 and Table 2.5 present cross-corpus synthesis results. They underlie the research gap and packaging design further applied.

2.9.1 TCCM summary

Table 2.4. TCCM synthesis — dominant findings and implications for this study.

Dimension	Dominant finding in the core	Implication for this study
Theory	Gravity / trade-area, geomarketing–GIS, MCDA + local spatial dependence, equal-weight composites (OECD, Opportunity Index), affordability scoring, expected-response allocation (Gupta & Steenburgh), inclusive-growth policy + constrained fair ads, festival calendar demand	Markets as spatial units, OI vs local momentum, readiness × affordability, EV ₹ shares, digital execution constraints, fairness as ruled constraint, festival features
Context	Mostly non-India settings in Scopus geomarketing (Germany suitability, Burkina broadband, St Petersburg districts, Bratislava food retail, Peruvian MSEs, Brazilian WebGIS), India policy + festival OA via other methods	Transfer is method and construct, not setting replication → motivate India public secondary panels
Characteristics	White-spot suitability, affordability scores, category × place demand, LBM usability, regional equity, moving-holiday / festive demand	Map to OI, GSDP/CPI proxies, propensity cells, digital caps, fairness floors, festival flags
Methodology	Bibliometrics, open geospatial MCDA/GWR, GIS + clustering, transactional prediction, WebGIS, policy synthesis, OA working papers / festival case studies	None is a full India FMCG secondary-panel Decision Stack with holdout propensity, EV shares, and ruled fairness

Gap-fill texts cover expected-response budgeting on their own. They also cover fairness and festival demand. One gap remains. To the best of the reviewed Scopus and OA corpus, no downloaded paper joins all three. None frames response-weighted ₹ shares, equity floors, and festival seasonality in one India geomarketing Decision Stack on public secondary panels.

2.9.2 ADO summary

Table 2.5. ADO synthesis — antecedents, decisions, and outcomes mapped to the Decision Stack.

ADO element	What the corpus emphasises	Analogue in this study
Antecedents	Spatial demand, competitive white spots, digital divide, shock/uncertainty, mobile ubiquity, regional dualism, festival calendar spikes	Panel features: CPI, population, clutter, TRAI readiness, festivals, affordability
Decisions	Where to site/invest, technology/channel mix, assortment keep-or-close, allocate to expected response, rebalance lagging regions	OI ranking, digital-share cap, category propensity, EV shares, fairness floors (b)
Outcomes	Opportunity maps, priority coverage plans, productivity/usability results, intended regional equity, festive demand lifts	Priority cells, ₹ shares, fairness deltas, and playbooks as auditable decision surfaces for planning

2.10 Research gap

The thematic reading and the TCCM / ADO coding leave a split map. Geomarketing and fit work answer where opportunity looks strong. Marketing-metrics, inclusive-growth, and festival literatures answer how spend should track response. They also answer why lagging places deserve a floor and when Indian demand spikes. Those answers rarely share a geomarketing paper. On an open India state × category panel they almost never sit together.

Table 2.6. What each theory strand establishes. What remains missing for India FMCG planners who must work from public data.

Theory strand	What the literature already establishes	What remains missing for India FMCG planners with public data
Geo-marketing	Markets are spatially interdependent, GIS suitability and opportunity mapping are mature	Few open Decision Stacks at India state × category that end in auditable ₹ shares
Media metrics	Budget should track expected response (Gupta & Steenburgh, Farris, Lilien, Tellis)	Lift assumptions are often hidden, geomarketing hits rarely close expected-value budget mathematics
Elasticity / affordability	Prices and incomes shape demand	Paid panels dominate, public-proxy affordability is under-used in spatial allocation
Digital readiness	Connectivity conditions what media and fulfilment can execute	Readiness is rarely coded as a hard digital-share cap inside an expected-value allocator
Inclusive growth	Policy floors for lagging regions are well argued (NITI 2018, NITI 2023 ADP), constrained fair advertising shows equity can bind with limited efficiency loss (Celis et al., 2019, Gelauff/Ali et al., 2020)	Fairness floors are almost absent from geomarketing results, planners need a ruled, inspectable reallocation on India state × category panels
Festival seasonality	Indian calendars produce predictable category spikes (NIPFP, festive OA papers)	Festival pressure is almost absent from geomarketing allocation models

Theory strand	What the literature already establishes	What remains missing for India FMCG planners with public data
Composite weighting	Equal weights are a standard transparent default when preference surveys/AHP are unavailable (OECD, 2008, Poledníková, 2017, Opportunity Index, 2023, Zhao et al., 2024), Tudor (2025) shows weight alternatives can leave surfaces highly correlated	Few India FMCG geomarketing stacks state an equal-weight OI and report weight/assumption sensitivity on public panels
Integrated packaging	Separate papers cover site choice, channel mix, equity, or festivals	Within the reviewed corpus, no joint Opportunity Index × growth propensity → EV ₹ shares + fairness + festival + playbooks stack on secondary panels

Peer-reviewed geomarketing can already name strong places. The weaker question is the budget question. It asks how much of a fixed national budget each India state–category cell should receive. That question matters once near-term growth propensity is in view. So do an explicit response-lift assumption, digital-execution caps, festival dates, and a fairness floor for thinner markets. All are taken seriously on public secondary totals.

Fairness and festival logic already exist next door. Opportunity scoring and EV ₹ shares are joined to those strands here inside one geomarketing Decision Stack at the India state × category × month level. The claim is applied packaging. No new Opportunity Index formula is offered. No unique welfare best point for fairness is claimed either.

2.11 Conceptual framework

Markets remain spatial units. Structural opportunity is evaluated separate from short-term momentum. Budget shares predict anticipated response. The India-specific standard of fairness and digital execution standards restrict the calculation. The playbooks articulate the outcome in auditable language. Figure 2.2 illustrates the mapping of public panels to the above for evaluators.

The familiar concepts of geomarketing and resource allocation are combined here into a single Market Decision Stack. Structural opportunity, growth propensity, priority stance, and fair budget shares become related terms in relation to public secondary aggregates. The twelve hypotheses (H1-H12) of Chapter 1 make these strategies into Decision Stack rules. Formulas are in Methodology. Test results are in Results.

2.11.1 Unit of analysis

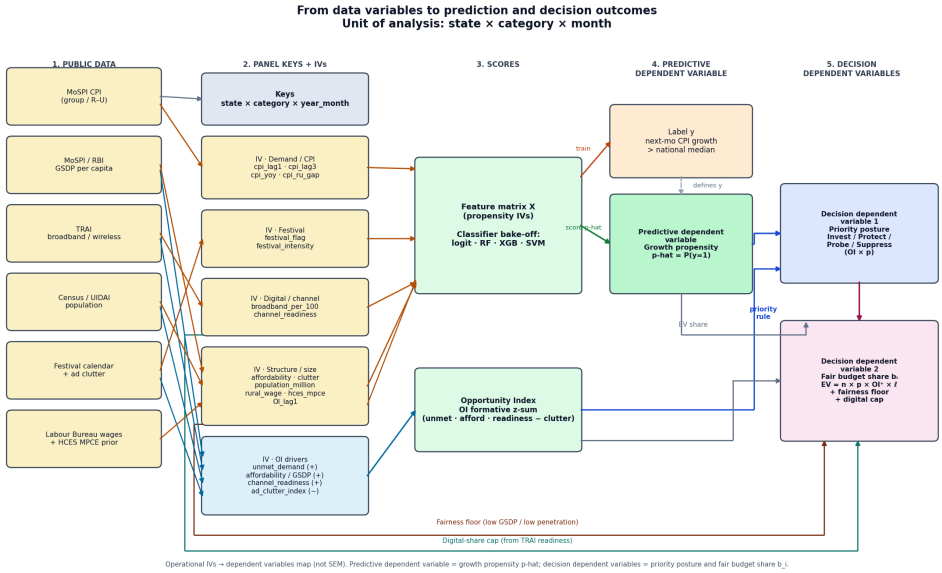
One of these observations is a cell of an India FMCG state × category × month. State and category are design keys defining the market. The lags, z-scores within months, and budget shares are calculated using the cell structure described.

The public layers have two purposes. For the Opportunity Index, unmet demand is population × category intensity versus a sales share estimate. Affordability is GDP per capita. Channel readiness is broadband/wireless density according to TRAI. Competitiveness is an advertising clutter index.

In the case of growth propensity, the list of features includes CPI persistence (lags, year-over-year change, rural-urban difference). It includes festival flags and strength. It includes lags in the Opportunity, population weight, and priors in safe-scope affordability such as rural wages and HCES-type MPCE shares.

The target in the case of growth propensity is the relative momentum. It is a next-month CPI growth above the national median on a month-to-month basis. After Opportunity Index and growth propensity scoring, the stack provides priority posture (Invest/Protect/Probe/Suppress). It also provides a fair budget share after EV allocation, equity floors, and digital ceilings.

From public data and panel IVs to predictive propensity and decision outcomes



From public data and panel IVs to predictive propensity and decision outcomes

Figure 2.2. Data-to-decision map. Public panels (CPI, GSDP, TRAI, population, festivals, clutter, wage/HCES priors) join on state × category × month keys. Those layers form the Opportunity Index and the propensity feature set. Growth propensity is the chance that next-month CPI growth beats the national median. Outcomes are priority stance (OI × propensity) and fair budget share (EV rule + fairness floor + digital cap).

2.11.2 Core constructs

The Spatial Opportunity (Opportunity Index) is an equal-weighted, clear composite within month. It consists of unmet demand, affordability, channel readiness, and ad clutter penalty. This is the opportunity concept discussed in Section 2.4. Equal weights are used in composite indicators without AHP elicitation (OECD, 2008; Opportunity Index, 2023). Scaling and the exact formula are described in Methodology.

Growth potential is the probability of next month's CPI growth above median. It is measured by the classifier calibrated on the hold-out sample. CPI lags, festival load, and digital readiness are among its inputs (Sections 2.6–2.8). The term is distinct from structural opportunity. Invest, Protect, Probe, and Suppress strategies can now be derived from the joint mix. Unconstrained ₹ weights respond to the EV rule. Floors and caps are set after (Sections 2.4 and 2.7).

2.11.3 How the strands join

Multimarket and fit research support analysis of state × category cells. They also support a structural opportunity score before any spend rule is applied. Gupta and Steenburgh's two-stage logic—estimate response, then allocate—gives the budgeting backbone. NITI-style lagging-region logic motivates ruled floors as open reallocation. Digital readiness studies motivate a hard digital-share cap. Festival calendars enter propensity features and an optional sensitivity filter.

Collapsing opportunity and propensity into one score would hide Protect versus Probe gaps. The framework keeps both axes until priority is set.

2.11.4 From framework to method

The setup is step-wise. Step one is panel join. Opportunity Index, propensity bake-off, priority matrix, EV ₹ shares, fairness and digital limits, and market playbooks follow. Sensitivity tables keep planning tests clear: a uniform opportunity shock, and a festival-cell filter. The framework ends

where the usable Decision Stack begins. It is one secondary-data path from public totals to open state–category budget advice.

CHAPTER 3
DATA

3.1 Principles

This is the only place where the labeling of the dataset takes place in this study. Other chapters (Methodology to Conclusion) will be carrying that label. However, no other data statement will be reiterated.

The operational dataset in the present case is a harmonized secondary panel. No-cost and citable official publications get appended to a state-by-category-by-month structure prior to generating the Opportunity Index and propensity scores. This labeling becomes relevant since the latter two might otherwise appear as native marketing panels. They are not. They are secondary analysis of published aggregates.

Table 3.1. Design rules that govern every data join in this study.

Rule	What it requires in practice
Secondary inputs only	Use free, citable public panels, do not substitute proprietary scanner or brand P&L extracts for the reported results
Traceability	Record every portal URL and retrieval date in Appendix-I, carry formal citations as References 37–48
Harmonisation before scoring	Map states and Union Territories through a canonical code table before CPI, GSDP, TRAI, or festival joins enter the Opportunity Index
Clear designation	Label all reported results as a harmonized secondary panel of published official statistics (coverage in Section 3.6, construction in Appendix-I)

These rules set the rest of the chapter: terminology (Section 3.2), source layers (Section 3.3), category map (Section 3.4), schema and coverage (Sections 3.5–3.6), reliability caveats (Section 3.7), demand-proxy definitions (Section 3.8), and ethics of secondary use (Section 3.9).

3.2 Terminology

Two labels are fixed here for the bind. Later chapters inherit them and do not redefine the dataset.

Table 3.2. Researcher terminology for the secondary panel used in this report.

Term used in this report	Meaning
Harmonized secondary panel	Multi-source secondary data assembled from published official statistics and joined to a common state × category × month unit of analysis
Published official statistics	Public MoSPI, RBI, TRAI, Census of India, NFHS/UIDAI, Labour Bureau, and related aggregate releases cited in References 37–48

In method terms, results are a secondary analysis of published totals on the Decision Stack schema. Models score relative structure and momentum on this panel. They do not build a primary FMCG sales ledger.

3.3 Data layers

There are nine layers in total for the public sector. Out of these, six layers contribute to the OI and propensities. The other three layers include the safe scope extensions which have affordability, category prior, and urban-rural pricing tension when the public value is there.

Table 3.3. Data layers, public sources, panel roles, and reference pointers.

Layer	Public source (cite)	Role in panel	Ref.
Demand / inflation	MoSPI CPI state / group tables	Category-linked price / demand pressure	37, 39
Affordability	MoSPI State Domestic Product, RBI Handbook of Statistics on Indian States	GSDP per capita → ability to pay	38, 40
Digital readiness	TRAI Performance Indicator Reports	Broadband / wireless density	41
Demography	Census of India, NFHS / UIDAI summaries	Population and urbanisation	42, 43, 44
Seasonality	National Portal of India calendar / gazette holidays	Festival flags and intensity	45
Competitive clutter	ASCI and published category ad-intensity aggregates	Clutter penalty in Opportunity Index	46
Rural wages (extension)	Labour Bureau Wage Rates in Rural India, RBI Handbook wage tables	Affordability blend IV	47, 40
HCES MPCE shares (extension)	MoSPI Household Consumption Expenditure Survey (2022–23 / 2023–24)	Structural category-share prior	48
CPI rural–urban gap (extension)	MoSPI CPI Rural / Urban / Combined (eSankhyiki)	Sectoral price-pressure feature	37, 39

Full URLs, retrieval freeze (2026-07-15 and extensions to 2026-07-26), variable lists, and coverage descriptions can be found in Appendix-I (Tables I.2-I.3). Portal references are References 37-48. Window is from 2022-01 to 2025.

3.3.1 Variable inclusion and exclusion

Examiners often ask why these layers and not literacy, road density, or e-commerce penetration. The rule was simple. Each layer had to fit a monthly state × category grain. It had to speak to unmet demand, affordability, channel readiness, clutter, seasonality, or propensity. It also had to come from free official totals, not a proprietary scanner. Table 3.4 records that filter.

Table 3.4. Included constructs versus commonly suggested alternatives that were set aside.

Decision	Construct / candidate	Why retained or set aside
Included	CPI (incl. rural–urban gap)	Monthly demand / price pressure on the decision grain
Included	Population, GSDP per capita	Size and ability-to-pay priors for OI and EV weights
Included	TRAI broadband / wireless	Public, state-linked channel readiness
Included	Festival flags / intensity	Indian calendar seasonality for propensity
Included	Ad-clutter index	Competitive penalty in the Opportunity Index

Decision	Construct / candidate	Why retained or set aside
Included (extension)	Rural wages, HCES MPCE shares	Affordability and category-share priors where published
In panel, not in OI equation	urban_share	Correlated with GSDP / TRAI readiness on this surface, kept for audit, not double-counted in OI
Excluded	Adult literacy	Thin monthly refresh, overlaps strongly with GSDP / urban structure for national FMCG media ranking
Excluded	Road / rail density	Infrastructure grain rarely matches monthly category planning, sparse comparable state series for the full panel window
Excluded	E-commerce GMV / app penetration	Mostly proprietary or irregular, TRAI density stands in as the public readiness proxy
Excluded	Brand offtake / scanner volume	Out of secondary-data scope (Chapter 1 non-goals)

The OI uses four scaled constructs: unmet demand, affordability, channel readiness, and clutter. Propensity adds CPI lags, festivals, readiness, lagged OI, and extension priors. Exclusions reflect data grain, not a claim that literacy or logistics lack commercial value.

3.4 Category taxonomy

Public CPI group labels do not match brand P&L trees one-for-one. This study maps public groups onto six FMCG decision categories shared by OI and propensity. Table 3.5 shows the map.

Table 3.5. Illustrative CPI / public-group map to six FMCG decision categories.

CPI / public group (illustrative map)	Decision category
Food & beverages (staples-heavy)	Food_Staples
Beverages / soft drinks related	Beverages
Packaged convenience foods	Packaged_Foods
Personal care	Personal_Care
Oral care	Oral_Care
Household care / cleaning	Household_Care

After state-name fixes, category mapping is the next big harmonisation risk. The map is checked before any claim that CPI groups match FMCG P&L lines. Appendix-I holds the build notes for that check.

3.5 Panel schema

Once layers and categories are fixed, the pipeline builds one row per state × category × month. Table 3.6 lists the core columns that survive into scoring and allocation.

Table 3.6. Processed panel schema — core columns and meanings.

Column	Meaning
state, state_code	Canonical geography
category	One of six FMCG decision categories
year_month	YYYY-MM
cpi_index	Category-linked price / demand pressure proxy
category_intensity	Relative category weight

Column	Meaning
sales_share_proxy	Observed relative activity proxy
population_million, urban_share	Size / urbanisation
gsdp_pc_inr_thousand	Affordability base
broadband_per_100, wireless_per_100	Digital readiness
cpi_rural, cpi_urban, cpi_ru_gap	Rural / urban CPI and gap (MoSPI-style extension)
rural_wage_inr	Rural wage (Labour Bureau–style extension)
hces_mpce_share	HCES category MPCE share prior
festival_flag, festival_intensity	Seasonality
ad_clutter_index	Competitive clutter
expected_demand, unmet_demand	Size × intensity vs observed gap
affordability, channel_readiness	Constructed feature layers

Keys in geography (state, state_code) and the category designation lock in the unit of analysis. The price, demography, digital, festival, and clutter variables power the Opportunity Index. Columns from extension data (cpi_ru_gap, rural_wage_inr, hces_mpce_share) add to the list of variables for the propensity features in Chapter 4. Variables like expected demand, unmet demand, affordability, and channel readiness are generated in the pipeline.

3.6 Panel coverage

The reported pipeline run yields the coverage counts in Table 3.7. Dataset type is stated in the same table so designation and size travel together.

Table 3.7. Panel coverage metrics from the reported pipeline run.

Metric	Value
Dataset type	Harmonized secondary panel (published official statistics, Appendix-I)
Panel rows	8,928
States	31
Categories	6
Months	48 (2022-01–2025-12, excludes 2020–2021)
Latest-month decision cells	186 (= 31 × 6)

Years 2020-2021 are excluded deliberately. The pandemic period demand and pricing behavior is not good for FMCG forecasting under normal conditions. Many publicly available time series become much noisier during that period. The stacking algorithm trains and distributes predictions for 2022-2025 years only.

Coverage of core components of CPI, population, and holiday joining scores is high. The extension components such as rural wages, HCES MPCE shares, and the CPI difference between urban and rural areas cover gaps where data is published. The propensity score training procedure imputes medians in case when the corresponding cell does not have any features.

3.7 Reliability

Panel reliability reporting tracks missing values on keys and continuous layers. On the current run, state, category, and year_month show 0% missing. CPI (with rural–urban gap), population, GSDP, TRAI densities, festivals, clutter, rural wages, and HCES shares are fully filled by construction. Joins and within-panel imputation close gaps before scoring.

That fill rate reflects the harmonised artefact, not raw portal dumps. Byte-identical portal downloads would bring back the frictions in Table 3.8. Those must be fixed before scores are trusted.

Table 3.8. Known frictions when raw portal downloads replace harmonised joins.

Friction	Why it matters for scoring
State name alias collisions	Duplicate or mismatched geographies unless names pass through the canonical state code table
CPI group coverage gaps	Incomplete Northeast / Union Territory series can thin category-linked demand pressure
TRAI LSA vs state mismatches	Licensed service areas do not always equal state boundaries used in the panel
GSDP vintage lag	Annual GSDP joined to a monthly frame lags the latest CPI months
Labour Bureau wage coverage	Rural wage series exist for a subset of states only
HCES wave vintage	MPCE shares are a cross-section prior, not a monthly series

3.8 Demand proxies

Unmet demand is defined as a structural gap. It is not a forecast residual from a sales model. Expected demand scales with market size and category weight (Eqn. 3.1). Unmet demand is the difference between that expectation and the observed activity proxy (Eqn. 3.2):

(expected_demand∝population×category_intensity (3.1))

(unmet_demand=expected_demand-sales_share_proxy (3.2))

Affordability and channel readiness come from GSDP per capita and TRAI broadband/wireless density. Both enter the OI as within-month z-scores so ranks stay scale-free across big and small states.

3.9 Ethics

The panel uses published official totals only. Those series are proxies for decision support. They do not replace primary brand sales, retailer offtake, or household microdata. Appendix-I shows how to cite the panel as a Decision Stack input so users keep the secondary nature of the evidence visible.

CHAPTER 4
METHODOLOGY

4.1 Research design

India FMCG Geo-Budgeting is executed via a two-step Market Decision Stack process here. The reasoning is a synthesis of predict-then-allocate from marketing resource allocation with multi-criteria opportunity score from geomarketing suitability literature.

Step one ranks. The Chapter 3 panel results feed an Opportunity Index structure and a time-based classifier of growth propensity. A cross-tab of the two yields Invest / Protect / Probe / Suppress cells.

Step two allocates. Shares of a fixed national budget according to EV are modified by India-specific fairness floors and digital share ceilings. Market strategy manuals and one consolidated research report come out of this run.

The market cell is state × category × month. Figure 2.2 places the concepts. Section 4.10 reviews the whole process and how Results metrics link to this unique reproducible run.

```
Public panels → state×category×month panel → Opportunity Index
→ Growth propensity
→ Priority cells → Expected-value shares → Fairness / digital constraints
→ Market playbooks → Compiled report
```

4.2 Data assembly

Chapter 3 and Appendix-I already fix panel construction, designation, coverage, and portal notes. Modelling stages take that joined state × category × month frame as given.

4.3 Opportunity Index

The Opportunity Index is a four-part composite kept deliberately transparent so boards can read each lever. Within each month, four constructs are standardised and combined with equal weights (Eqn. 4.1):

$$OI = z(\text{unmet_demand}) + z(\text{affordability}) + z(\text{channel_readiness}) - z(\text{ad_clutter}) \qquad (4.1)$$

Equal weights are a matter of design, not welfare optimality. In the absence of a board survey, equal weights are typically adopted as the obvious default choice in composite index research (OECD, 2008; Poledníková, 2017). Opportunity Index applications do likewise (Opportunity Index Technical Supplement, 2023). Equal weights have been used as a neutral starting point in GIS multi-criteria applications (Zhao, van Duynhoven & Dragićević, 2024). Tudor (2025) demonstrates that other weighting methods could result in extremely high correlation between OI surfaces.

There was no board AHP in this study. Equal z-weights within each month maintain the openness of OI to auditing. Weight and shock tests will be carried out in Section 4.8. Alternative weighting methods have been dropped. PCA would use variance as importance in the absence of board input. Entropy would overweight z-component noise rather than the four levers. Regression weighting needs brand volume or media ROI, which this sample lacks. Chapter 6 stresses test the robustness of an affordability-based weighting against the equal weights surface.

Unmet demand = Expected demand (Population * Category intensity) – Sales share proxy. Affordability is a monotonic transformation of per capita GSDP. Channel readiness is calculated using broadband density of TRAI and wireless density. Advertising clutter is included with a minus

sign so that cluttered cells are penalized. High Opportunity Index score indicates the presence of residual demand, ability to pay, and execution readiness. As scores represent z-sums for the month, they indicate structural attractiveness relative to the cross-section.

4.4 Growth propensity

Prediction target

Next-month growth in the state x category CPI is compared to the national median of that month, and the binary label will be one when the growth is higher than that median. Relative momentum is thus the target. The level of absolute inflation is not part of the label. A simple positive growth flag is not a part of the label.

The classification task works well within the stack hierarchy since the prediction output is the priority posture (Invest/Protect/Probe/Suppress). Even if the demand were a continuous variable that requires ranking of the budgets, the binary label based on the relative median incorporates the ranking logic into the holdout ROC-AUC/AP. Chapter 6 validates the decision by comparison with Ridge-based CPI growth model on the same holdout. Oftake validation awaits proprietary data series.

Features

The feature set mixes demand persistence, seasonality, digital capacity, and structural context.

Table 4.1. Propensity-model feature set and roles.

Feature	Role
cpi_lag1, cpi_lag3, cpi_yoy	Price / demand dynamics
festival_flag, festival_intensity	Seasonality
broadband_per_100, channel_readiness	Digital / channel capacity
OI_lag1	Lagged opportunity
affordability, ad_clutter_index, population_million	Structure / size
cpi_ru_gap	Rural–urban CPI differential (safe-scope extension)
rural_wage_inr	Rural wage affordability prior (safe-scope extension)
hces_mpce_share	Category MPCE share prior (safe-scope extension)

CPI lags are pre-specified classifier features chosen for planning cadence. Lag-1 captures short-run persistence. Lag-3 captures quarter and pre-festival memory. Year-on-year change captures the seasonal horizon.

Table 4.2. Pre-specified CPI lag features — construction and role.

Feature	Construction	Role
cpi_lag1	CPI level, (t-1)	Short-run persistence
cpi_lag3	CPI level, (t-3)	Quarter / pre-festival memory
cpi_yoy	12-month % change	Year-on-year + seasonal horizon

Nested holdout diagnostics and stationarity notes for this lag design are reported in Appendix-I (Section I.7).

Holdout and model comparison

Months are sorted chronologically. The last three months form the evaluation holdout. Earlier months train the models. Four classifiers are compared on identical features and the same holdout:

1. logistic regression (L2, standardised features, class_weight="balanced")

2. Random Forest (200 trees, max depth 10, min leaf 5, balanced)
3. XGBoost when native OpenMP support is available, otherwise an sklearn gradient-boosting fallback labelled honestly in results
4. a linear SVM with sigmoid calibration for probabilities

The binary label is relative (next-month CPI growth above the national median). The positive rate is about 0.5 by definition. Classical rare-event imbalance is not the concern. Balanced class weights are used, however, as an extra precaution and for consistency in the bake-off.

Model selection is done using the maximum ROC-AUC on the holdout data, with average precision as tiebreaker. Logistic regression is chosen in case of residual tie as the model is simple and coefficient estimates are interpretable (standardised coefficients that can be audited for board presentation). Model probabilities become growth_propensity for decision-making. Scoring is done for the panel that we consider usable. The latest month's cell snapshot transfers propensity to priority and budget modules.

Linear logit model wins the bake-off on this panel (Chapter 6). Feature space is limited. The label is relatively balanced. Tree methods do not bring much value after CPI year-on-year and affordability. Linear logit model stays because of high holdout performance and interpretability of the coefficients. Uniqueness of the learner is not claimed. Six-month holdout, walk-forward windows and 2024→2025 data split are present in Results and Appendix-I.8.

4.5 Priority cells

Having opportunity index and growth propensity on distinct axes is intentional. Cells can be structurally attractive but weakly momentum-driven; cells can be momentum-strong but structurally hollow. Median splits of the two indicators on the latest-month surface give rise to four manager mindsets.

Median splits are default because they are unambiguous. They provide roughly equal numbers of cells for planning meetings. No assumption about the number of clusters is required. Quartile or 60th-percentile splits would reduce the size of the Invest cell set. The use of k-means or Jenks classifications would create arbitrary boundaries when there is no brand volume to determine "true" segments. The agreement of the median splits and 60th-percentile splits is reported in Chapter 6 as a robustness check. Absolute business thresholds (such as a minimum of 0.55) are still available to boards preferring cutoffs over splits.

Table 4.3. Priority cells from median splits of Opportunity Index × growth propensity.

	High growth propensity	Low growth propensity
High OI	Invest	Protect
Low OI	Probe	Suppress

Invest cells combine high opportunity with high propensity. Protect markets stay structurally strong when near-term momentum cools. Where momentum outruns the Opportunity Index, Probe is the planning label. Suppress is reserved for cells that score low on both structural opportunity and near-term growth propensity. Campaign language for each posture is written into automated market dossiers.

4.6 Expected-value shares

Unconstrained shares follow an expected-value form so that budget tracks response potential. For each latest-month cell (i), the positive OI transform (Eqn. 4.2), expected value (Eqn. 4.3), and normalised share (Eqn. 4.4) are:

$$OI_i^+ = OI_i - \min_j OI_j + \varepsilon \qquad (4.2)$$

$$EV_i = n_i \times p_i \times OI_i^+ \times \ell \qquad (4.3)$$

$$b_i^{EV} = B \times \frac{EV_i}{\sum_j EV_j} \qquad (4.4)$$

Here (n_i) is population in millions (floored), (p_i) is growth propensity, (OI^{+}_i) is a positive transform of the Opportunity Index, (B) is the illustrative national budget of ₹10 crore, and (= 0.04) is an explicit planning assumption for response lift. The board can sensitivity-test the parameter ().

The product formula was used deliberately since size, propensity, and positive OI all have to be taken into account. If we use an additive score, the most influential variable may overshadow other variables. If we use a log form, the differences we want to see from the perspective of a ₹10 crore budget will not be revealed anymore.

Shares are normalized by (_j EV_j). Changing () scales total EV but not relative shares under this formula. Lift is a stated planning assumption about the size of the response. The rank order does not change with () using the above formula. A 4% lift is a conservative planning prior in the absence of brand media test data. This panel is not suited to estimate (). Boards with internal lift data may change ().

4.7 Fairness and digital caps

The expected-value shares are then subjected to constraint ruling adjustments. National plans remain visibly present in thinner markets. Digital weight corresponds to TRAI readiness.

Fairness-eligibility occurs in the case where the per-capita GSDP is either below or at the 35th percentile, or penetration (sales divided by expected demand) is below or at the 35th percentile. Every eligible cell receives a floor. The overall floor pool is limited to 40% of B. Every cell's budget equals the greater value between EV share and floor. The shares are renormalized such that they add up to B. The cells with low TRAI readiness are subject to a 25% cap on digital share assumption. Invest and Probe are more digital than Protect and Suppress. The fairness delta records the rupee change from before to after constraints.

These thresholds become parameters in the discussion among board members. The rationale derives from the logic of lagging regions focus (NITI Aayog), and allocation in terms of equity and inspectability (Celis et al., 2019, Gelauff et al., 2020). Stated thresholds represent declared policy choices and not optimal choices by survey respondents. The 35th percentile thresholds flag about one-third of cells as eligible. This range is sufficiently large enough to be relevant while being limited to a degree that floors don't devour the entire budget. 40% pool threshold becomes a limit to the change in ranks of EVs. 25% digital threshold provides some digital capacity but without digital-first policy in broadband-scarce states. These thresholds become board decisions in practice. Chapter 6 tests floor pools at 20% / 40% / 60%. Changes in digital thresholds impact only channel mix. Total cell ₹ remains unchanged on this simulation.

4.8 Sensitivity

Planning scenarios show how the surface moves when assumptions change. Index QA treats weight and assumption tests as standard practice (OECD, 2008). Uniform OI shocks and festival-only filters can leave Invest totals flat when the min-shift transform is rank-invariant or when every latest cell is festival-flagged. Useful checks stress category OI shocks, OI weight tilts, and fairness floor-pool size.

Table 4.4. Sensitivity scenarios — construction and managerial reading.

Scenario	Construction	Managerial reading
Baseline	Full stack on latest cells	Planning default
Uniform OI +0.5	Add +0.5 to every cell's OI before EV	Rank-invariant under min-shift, diagnostic only
Food_Staples OI +0.5	Shock OI only in Food_Staples	Category-specific pressure

Scenario	Construction	Managerial reading
Affordability-heavy OI	Weights 0.5 / 1.5 / 1 / 1 on the four z-components	Board tilt toward ability to pay
Fairness floor pool 20% / 40% / 60%	Cap collective floors at that share of (B)	How hard equity binds
Digital cap 15% / 25% / 40%	Annotate assumed digital share in low-readiness cells	Channel-mix policy (₹ totals unchanged on this run)

Because (OI^{+}) is min-shifted, adding the same constant to every cell leaves relative expected-value ranks unchanged. Category- or weight-specific shocks, and floor-pool alternatives, are the informative path reported in Results.

4.9 Design decisions summary

Table 4.5. Method choices retained for this study, with the main alternative set aside.

Choice	Retained	Main alternative set aside	Warrant
Unit of analysis	State × category × month	District / city / store	Public MoSPI–TRAI–GSDP series align cleanly at state, finer grains lack consistent monthly category panels here
Opportunity weights	Equal within-month z-weights	PCA / entropy / AHP / regression weights	No board AHP, OECD-style transparent default, tilt sensitivity reported
Propensity task	Relative-median classification	Continuous CPI / volume regression	Priority postures need ranks, AUC/AP match that job, Ridge triangulation in Results
Propensity learner	Logistic regression (bake-off winner)	RF / boosting / SVM	Holdout AUC/AP + coefficient auditability, stable across 3-/6-month and walk-forward checks
Priority cut	Median split	Quartiles / k-means / business threshold	Transparent 2×2, p60 agreement checked
Budget form	Multiplicative EV	Additive or log EV	Size × momentum × opportunity must co-occur, () disclosed
Fairness / digital	35th pct eligibility, 40% floor pool, 25% digital cap	Estimated social optimum	Policy levers under NITI-style lagging-region logic

4.10 Outputs

The Decision Stack produces three classes of output from one run:

1. diagnostic figures for the propensity bake-off, Opportunity Index bands, priority cells, and fairness adjustments.
2. market playbooks that state each cell's priority posture, scores, and recommended ₹ share.
3. a compiled research report that binds the chapters, live analytics tables, and playbooks.

Public panels aggregate into state × category × month framework. Opportunity indices and growth propensities are calculated. Priority cells and EV shares apply fairness and digit rules. Graphs and playbooks are derived from these calculations. The report is assembled for inspection. Holdout ROC-AUC, mean precision, fairness metadata, and budgets from this run serve as the live source for Abstract and Results values.

Repeat the identical module pipeline to regenerate the surface layer: ingest → panel → opportunity → propensity → priority → allocation → graphs → playbooks → assemble. Using fixed seeds, the holdout criteria from Section 4.4, and the formulae from Sections 4.3 through 4.7, the results are auditable. Defence diagnostics (python scripts/defence_diagnostics.py) regenerates OI–propensity associations, weight tilt robustness, split priority consensus, and sensitivity tables in data/processed/.

CHAPTER 5
NOVELTY AND CONTRIBUTION

5.1 Contribution

There are existing suitability scores, expected response budgets, off-the-shelf classifiers, and floors for inclusive growth. However, each one tends to run on its own. This is an example of applied packaging. They have been combined into one Market Decision Stack panel for India state × category markets. Budget allocation and market playbooks depart from the same starting point.

The existence of any Opportunity Index formula, unique classifier family, causal media-mix return on investment, or social optimum is not claimed here.

5.2 What this study adds

Brand sales ledgers sit outside public panels. Test-grade response curves do too. Inside that limit, Opportunity Index and growth propensity stay on separate axes. The joint priority surface becomes EV ₹ shares under lift (= 0.04). India-specific fairness floors and digital-share caps bind next. Playbooks ship with the compiled report. Formulae live in Methodology. Live tables live in Results.

5.3 Relation to prior work

Three anchors position the packaging statement.

Suitability / white-spot geomarketing (for instance, Tudor, 2025) identifies high-potential regions. In such cases, the ranks become portions of an established national marketing budget after adding Indian fairness considerations and digital caps.

Expected-response allocation (Gupta & Steenburgh, 2008) holds that spending must align with anticipated response. With this board, panel lift remains a predetermined planning assumption. The boards may monitor the ₹ table with () varying.

Inclusive-growth guidance from NITI Aayog on low-achieving regions leads to minimums for thinly populated markets. These minimums rest on the OI x propensity matrix. Equity remains auditable alongside growth.

The connection is the contribution. The axes remain separate. The portions remain clear. The rules remain open. The playbooks emerge from one iteration.

5.4 Evidence map

Contribution	Where it is shown
Two-axis priority (OI × propensity)	Results and findings: priority tables
Holdout-selected growth propensity	Results and findings: propensity bake-off
EV shares + fairness / digital rules	Results and findings: allocation and fairness metadata
Stress-test reporting	Results and findings: sensitivity tables
Market playbooks from the same run	Bound market playbooks in the compiled report

5.5 Planning use

Table 5.1 is the meeting surface. Cadence and stance sit in Managerial implications.

Table 5.1. What the Decision Stack delivers for planning.

Planning need	What this study delivers
Cell ranking with actions	Invest / Protect / Probe / Suppress postures
₹ shares of a fixed budget	EV shares plus fairness deltas
Separate structure and momentum	OI and propensity on two axes
Inclusive geographic reach	Fairness floors with stated cutoffs
Operating artefacts	Market playbooks bound with the report

CHAPTER 6
RESULTS AND FINDINGS

Observed results for the Decision Stack exercise on the Chapter 3 panel are shown below. This chapter discusses topics such as research scale, propensity bake-off, Opportunity Index range, prioritization counts, budget constraints, sensitivity tests, and the relationship between the hypothesis H1-H12 and the data. Analysis follows the tables very closely. Media ROI causality is not included here.

6.1 Study scale

Table 6.1 summarises the analytic frame. The working panel contained 8,928 rows: 31 states × 6 FMCG categories × 48 months within 2022-01–2025-12. Budget shares were computed on the latest-month surface only (186 cells). The illustrative national budget was ₹10 crore. Response lift was fixed at (= 0.04) for the EV step. Propensity models were scored on a 3-month chronological holdout. These quantities fix the scale for every result that follows.

Table 6.1. Pipeline scale — panel size, budget, assumed lift, and holdout months.

Metric	Value
Panel rows	8,928
States × categories × months	31 × 6 × 48
Panel window	2022-01 – 2025-12
Latest-month decision cells	186
Illustrative national budget (B)	₹10,00,00,000
Assumed response lift ()	0.04
Holdout months	3

6.2 Propensity model comparison

Identical characteristics and same hold-out period were possessed by four classifiers. Winner was selected based on highest hold-out ROC-AUC, with average precision settling ties. Logistic regression solved remaining ties. Positive class was defined to be CPI growth greater than national median in next month. Hold-out positive rate became 0.5 due to construction. Problem of rare event did not become a primary concern. class_weight="balanced" remained in effect throughout bake-off as insurance. Discrimination in this case is relative momentum across markets.

Table 6.2 reports the bake-off. Logistic regression reached holdout ROC-AUC 0.761 and average precision 0.752. Calibrated linear SVM was a hair behind (0.760 / 0.750). Tree ensembles trailed further: XGBoost/GBM fallback 0.739 / 0.726, and Random Forest 0.736 / 0.724. Table 6.3 and Figure 6.1 record the selected model on 8,184 training and 372 holdout rows. Scores above chance (0.5) entered later modules as growth_propensity.

Table 6.2. Propensity model bake-off — holdout ROC-AUC, average precision, and selection.

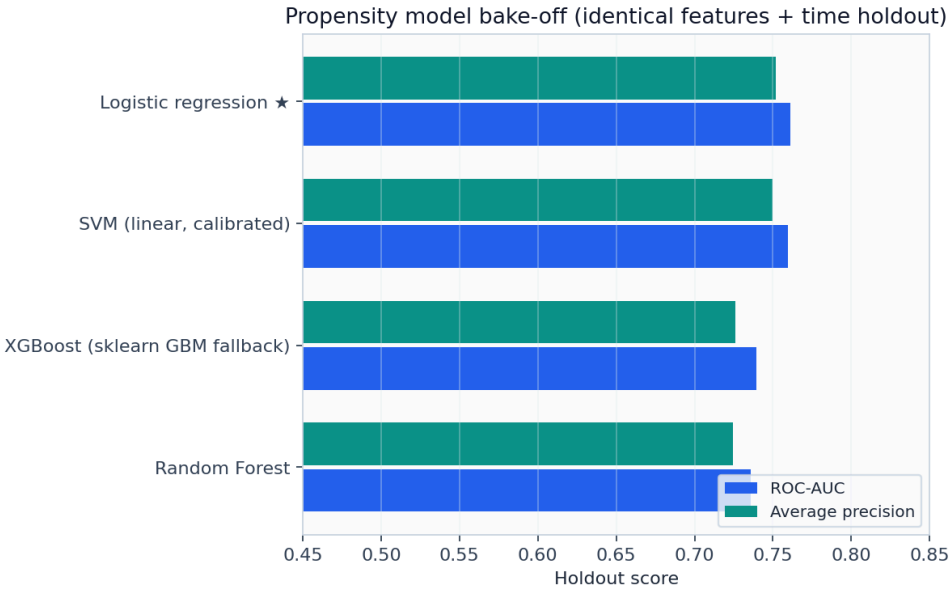
Model	Holdout ROC-AUC	Holdout AP	Selected
Logistic regression	0.761	0.752	★
SVM (linear, calibrated)	0.760	0.750	
XGBoost (sklearn GBM fallback)*	0.739	0.726	
Random Forest	0.736	0.724	

*On the reported run, native XGBoost OpenMP support was unavailable, so a sklearn gradient-boosting classifier stood in for the boosting family. A machine with native XGBoost may move that row slightly. The selection rule stays the same.

Table 6.3. Selected propensity model metrics (logistic regression).

Metric (selected model)	Value
Selected model	Logistic regression
Train rows	8,184
Holdout rows	372
ROC-AUC	0.761
Average precision	0.752

Propensity model bake-off holdout ROC-AUC and average precision



Propensity model bake-off holdout ROC-AUC and average precision

Fig. 6.1. Holdout bake-off under identical features and time holdout.

Selected model coefficients

Table 6.4 lists the largest absolute standardised coefficients from the retained logistic fit (training-window standardisation, class_weight="balanced"). Signs are one-way links with the above-median growth label on this public panel. Year-on-year CPI carried the strongest weight (−0.75). That pattern fits mean-reversion against a relative-median label. Affordability entered positively (+0.31). Recent CPI level (cpi_lag1) was negative (−0.28). Rural wages, lag-3 CPI, HCES MPCE share, channel readiness, and lagged OI were smaller in magnitude. Festival flag and intensity were near zero once CPI and affordability were controlled.

H5-H7 view festival pressure, demand persistence, and digital preparedness as attributes expected to improve propensity. Holdout AUC/AP validate the combined attribute set as a ranker. Individual indicators showed mixed results. Affordability and preparedness favored positive contribution. CPI dynamics dominated with negative relative median correlations. The influence of festivals was close

to zero for this linear model. This research considers H5-H7 as design ideas, while holdout AUC/AP is the primary empirical validation of the propensity model. Nested lag models, ACF/PACF, and stationarity remarks are available in Appendix-I (I.7).

Table 6.4. Largest standardised logistic coefficients (absolute size).

Feature	Coefficient	Plain reading
cpi_yoy	−0.75	Strong year-on-year CPI change associates with lower odds of beating the national median next month (mean-reversion / relative ranking effect)
affordability	+0.31	Higher GSDP-linked affordability associates with higher propensity
cpi_lag1	−0.28	High recent CPI level associates with lower relative next-month odds
rural_wage_inr	−0.08	Small negative association after other controls
cpi_lag3	−0.08	Mild quarter-lag association
hces_mpce_share	+0.06	Slightly higher propensity where category MPCE share priors are larger
channel_readiness	+0.04	Mild positive link to readiness
OI_lag1	+0.04	Lagged opportunity enters with a small positive weight
festival_flag / festival_intensity	≈ 0	Festival terms are small once CPI and affordability are in the model

Cross-checks — continuous baseline, calibration, holdout stability

Three internal checks test whether classification, the kept logit, and the three-month holdout are fragile design choices.

Table 6.5. Continuous Ridge baseline versus logistic propensity on the same features and 3-month holdout.

Model	Level fit	Ranking score (vs above-median label)
Ridge on next-month CPI growth	RMSE ≈ 0.009, $R^2 \approx -1.59$	AUC 0.603, AP 0.610
Logistic classifier (same holdout)	—	AUC 0.759, AP 0.748

The Spearman correlation between Ridge scores and logit odds was 0.415. Top quartile agreement is around 48%. The continuous-level forecast is poor in this public panel. It ranks below the relative-median classifier. This reinforces the need to maintain the classification layer as the propensity layer for priority postures.

Holdout logit odds had a Brier score of 0.207. Bins increased with odds predicted. The middle bins slightly overestimated positives. The absolute odds are handled cautiously. Ranks continue to supply expected value proportions.

Table 6.6. Holdout-length bake-off and walk-forward stability (winner by ROC-AUC).

Check	Winner	Logit ROC-AUC
Last 3 months (main bake-off)	Logistic regression	0.761
Last 6 months	Logistic regression	0.749

Check	Winner	Logit ROC-AUC
Walk-forward ends 2025-08 … 2025-11 (vs RF)	Logistic regression (all four)	0.726–0.741
Train ≤2024-12 / test 2025 (calendar split)	Logistic regression	0.734 (AP 0.729)

The calendar split is still a panel-level check. It extends the time holdout. It is not an industrial sell-out back-test. District-grain and brand-volume external checks stay out of scope (Chapter 8, Appendix-I.8).

6.3 Opportunity Index

OI values are within-month z-sums. The panel mean hovers around zero by construction. As per the most recent month's surface, the four priority bands have been differentiated on structure. Invest cells yielded OI ≈ 1.39 (with a range of ≈ −0.3 to 5.4). Protect had a relatively high score on structure (≈ 1.86) but low on propensity. Probe had a low structural score (≈ −1.14), but close to the median propensity. Suppress had the lowest OI (≈ −1.88). Keeping OI and propensity separate is how we make the four-dimensional matrix meaningful.

The axis-dependency test was done using the same most-recent-month surface (186 cells). Pearson correlation between OI and growth propensity was 0.086 (Spearman ≈ 0.199). There is very little relationship between the two in aggregate. This is consistent with H8's two-axis framework. Category splits exist. Food_Staples exhibited a clear negative relationship (Pearson ≈ −0.73). Packaged_Foods and Personal_Care demonstrated a moderate positive relationship. Protect versus Probe differentiations require careful interpretation in case of staples, where structure and momentum may differ greatly. Component VIFs for the four OI z-scores remain below 5 (the highest is for channel readiness ≈ 4.45). The equal-weight sum is not a near-collinear mash of identical layers.

Equal versus affordability-heavy OI weights (0.5 / 1.5 / 1 / 1 on unmet demand, affordability, readiness, clutter) remained highly aligned (Pearson ≈ 0.948, Spearman ≈ 0.922). That echoes Tudor-style evidence that weight tilts need not redraw the whole opportunity map. The tilt still moves budget mass in Section 6.9 because expected value multiplies the rescaled OI.

6.7 Priority matrix

Table 6.7 reports cell counts and band means on the latest-month surface. Invest and Suppress each held 55 cells. Median splits on two axes tend to balance counts. Probe and Protect held 38 each. Mean propensity differed only modestly across bands (about 0.51–0.54). Mean OI differed sharply. Population was highest in Suppress and Probe (56.4 and 44.3 million on average). Large, structurally weak markets still matter once population weight enters expected value.

Distinct OI means across bands fit the two-axis design (H8). Invest cells combined high OI with above-median propensity. That matches the Invest rule behind H9. A stricter 60th-percentile gate on both axes agreed with median labels on about 81% of cells. It cut the Invest set from 55 to 37. Median splits remain the more inclusive planning default. They are also stable under a tighter cut.

Table 6.7. Latest-month priority matrix — cell counts, mean OI, mean propensity, mean population.

PriorityCell	n_cells	mean_OI	mean_propensity	mean_population (mn)
Invest	55	1.39	0.544	27.0
Probe	38	−1.14	0.541	44.3
Protect	38	1.86	0.510	25.8
Suppress	55	−1.88	0.514	56.4

6.8 Budget allocation

After expected-value shares and fairness / digital rules, Table 6.8 shows Invest receiving ₹3.23 crore (32.3%) of the illustrative ₹10 crore budget. Probe received 22.6%, Protect 21.7%, and Suppress 23.5%.

Table 6.8. Fairness-constrained recommended budget by priority cell (₹ crore and share of B).

PriorityCell	n_cells	Recommended budget (₹ crore)	Share of (B)
Invest	55	3.23	32.3%
Probe	38	2.26	22.6%
Protect	38	2.17	21.7%
Suppress	55	2.35	23.5%

Probe and Suppress continued to comprise approximately one-fifth each of the budget. Expected Value is calculated as the product of Population weight, Propensity, and Positive Opportunity. Probe and Suppress contained substantial populations for this snapshot (around 3,105 and 1,684 million person-category units respectively). Despite having low OI, Size kept EV massive even before floors were applied. Fairness thus moved yet another rupee layer to qualified thin markets. The ranking, weighing, and floor presence was maintained even for weaker cells.

Table 6.9 records the fairness layer. 79 cells met the low-GSDP or low-penetration rule against a ₹4.0 crore floor pool (≈ ₹5.06 lakh per eligible cell). Positive fairness reallocation was about ₹1.44 crore. Eligible cells together held roughly half of post-fairness budget on this run. Digital share was capped at 25% where TRAI readiness fell below the broadband cutoff.

Table 6.9. Fairness constraint metadata — eligible cells, floors, cutoffs, and reallocation.

Item	Value
Fairness-eligible cells	79
Floor pool	₹4.0 crore
Floor per eligible cell	≈ ₹5.06 lakh
Positive fairness reallocation	≈ ₹1.44 crore
GSDP cutoff (35th pct)	135.7 (₹ thousand pc)
Broadband cutoff (35th pct)	32.98 per 100

A single cell demonstrates the dual level of logic. Uttar Pradesh × Food_Staples was an Invest cell (OI ≈ 1.27, propensity ≈ 0.54). The expected-value logic allocated a very high pre-fairness share (about ₹50 lakh). Post-redistribution, the suggested budget was about ₹42.2 lakh (fairness gap ≈ −₹7.8 lakh). Large Invest cells could fund floors for others without affecting their ranking. Shares were unconstrained as per expected value (H10). Floors increased the proportion of lagging cells (H11). The low readiness digital ceiling capped in areas of poor broadband (H12). All three worked as intended.

6.9 Sensitivity

Varied stress scenarios replace the earlier uniform-only table. Table 6.10 reports Invest budget and positive fairness reallocation under baseline, a Food_Staples-only OI shock, an affordability-heavy OI weight tilt, fairness floor pools at 20% / 40% / 60% of (B), and digital-cap annotations at 15% / 25% / 40%. A uniform +0.5 OI shock remains rank-invariant after the positive transform. It is reported only as a formula check (Invest total unchanged at ≈ ₹3.23 crore).

Table 6.10. Discriminative sensitivity — Invest budget and fairness-pool delta by scenario.

Scenario	n_cells	Invest budget (₹ crore)	Fairness-pool delta (₹ crore)
Baseline (40% floor pool, 25% digital cap)	186	3.23	1.44
Uniform OI +0.5	186	3.23	1.44
Food_Staples OI +0.5	186	3.25	1.45
Affordability-heavy OI weights	186	2.89	1.50
Fairness floor pool 20%	186	3.38	0.59
Fairness floor pool 40%	186	3.23	1.44
Fairness floor pool 60%	186	3.06	2.20
Digital cap 15% / 25% / 40%	186	3.23	1.44

Floor-pool choice is the strongest lever on this surface. Tightening floors to 20% of (B) raises Invest mass to ≈ ₹3.38 crore and cuts positive fairness reallocation to ≈ ₹0.59 crore. Loosening to 60% does the reverse (Invest ≈ ₹3.06 crore, fairness delta ≈ ₹2.20 crore). The affordability-heavy OI tilt reduces Invest to ≈ ₹2.89 crore even though the tilted index stays highly correlated with the equal-weight OI. Budget shares feel weight policy more than rank correlations alone suggest. Digital-cap alternatives annotate assumed channel mix inside low-readiness cells. They do not reshuffle total cell ₹ on the current allocator. Festival-only filtering was not informative on the 2025-12 snapshot because every latest-month cell carried festival_flag = 1. Coverage by month is logged in Appendix-I.

6.11 Hypothesis outcomes

Table 6.11 maps design hypotheses H1–H12 to their role and to the evidence reported above. Holdout AUC/AP is the main stats check for propensity (H5–H7). OI band separation and the Invest rule speak to H1–H4, H8, and H9. Budget, fairness, and digital-cap tables implement H10–H12 as transparent allocation rules on the public panel.

Table 6.11. Design hypotheses H1–H12 — how this study treats each claim.

Hypothesis	Role in this study	Evidence in this chapter
H1–H4 (OI drivers)	Built into the OI formula	OI bands differ sharply across priority cells (§6.3–6.7)
H5–H7 (propensity features)	Features in the classifier	Holdout AUC/AP, mixed individual logit signs (§6.2)
H8 (two axes)	Design rule	Distinct OI means, similar propensity means (§6.7)
H9 (Invest posture)	Median-split matrix	55 Invest cells with high OI + higher propensity (§6.7)
H10 (EV shares)	Allocator formula	Budget table before/after fairness (§6.8)
H11 (fairness floors)	Ruled reallocation	79 eligible cells, ≈ ₹1.44 crore positive deltas (§6.8)
H12 (digital cap)	Hard cap at 25%	Applied on low-readiness cells (§6.8)

6.12 Evidence checklist

Table 6.12 closes the chapter with a compact evidence map for examiners and for any manuscript extract of these results. It lists secondary-panel sources, bake-off artefacts, selected propensity, priority and budget tables, sensitivity, and lag diagnostics.

Table 6.12. Results evidence map.

Evidence block	Where it appears in this report
Secondary panels	Harmonized CPI (R/U), GSDP, TRAI, population, festival, clutter, rural wages, HCES MPCE (2022–2025), portals in Appendix-I
Model bake-off	Fig. 6.1, Tables 6.2–6.3
Selected propensity	Logistic regression
Priority / budget	Latest-month cells, Tables 6.7–6.9
Sensitivity	Table 6.10 (discriminative), sensitivity_discriminative.csv
Lag diagnostics	Appendix-I (Section I.7)
Axis / weight / split checks	Appendix-I (Section I.8), OI–propensity corr, weight tilt, median vs p60
Tier-2 cross-checks	Tables 6.5–6.6, continuous vs classifier, calibration, rolling / calendar holdout

Analytic figures

These charts accompany Results and findings. Lag-design diagnostics appear in Appendix-I (Section I.7).

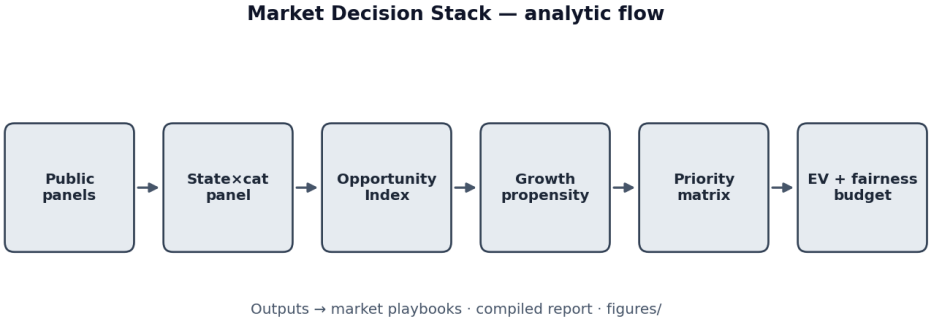


Fig. 6.2. Market Decision Stack analytic flow

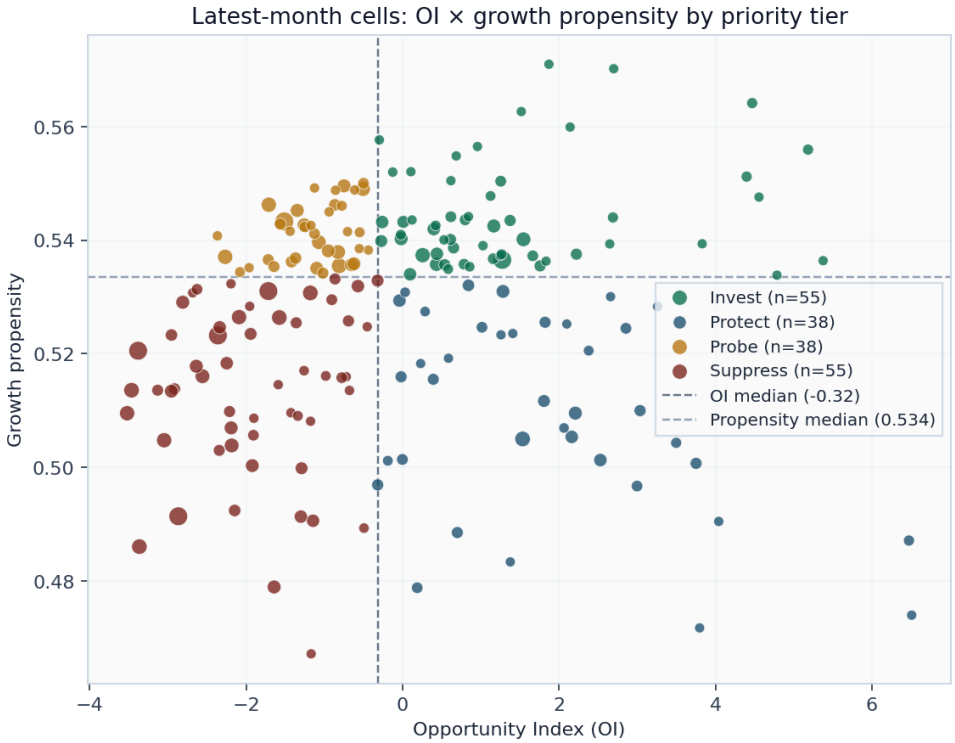


Fig. 6.3. OI × growth propensity by priority tier

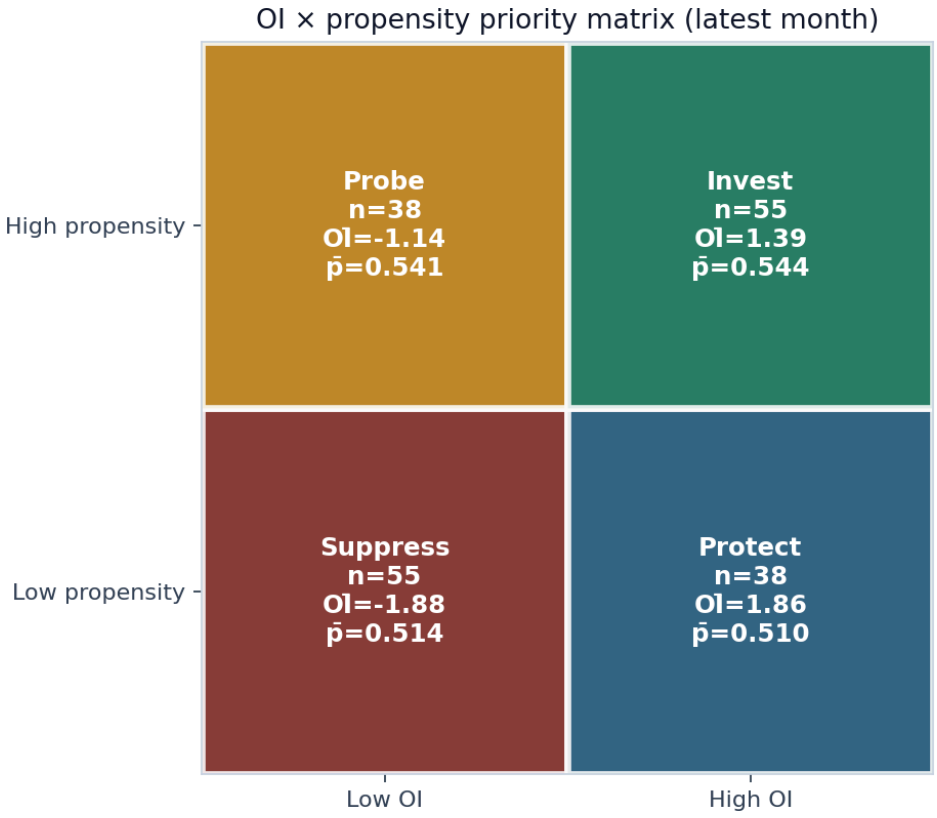


Fig. 6.4. OI × propensity priority matrix (latest month)

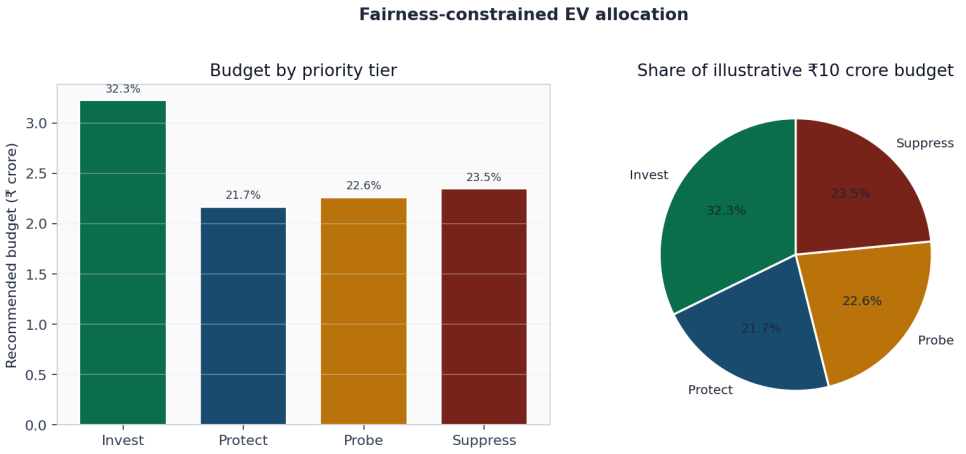


Fig. 6.5. Fairness-constrained budget by priority tier

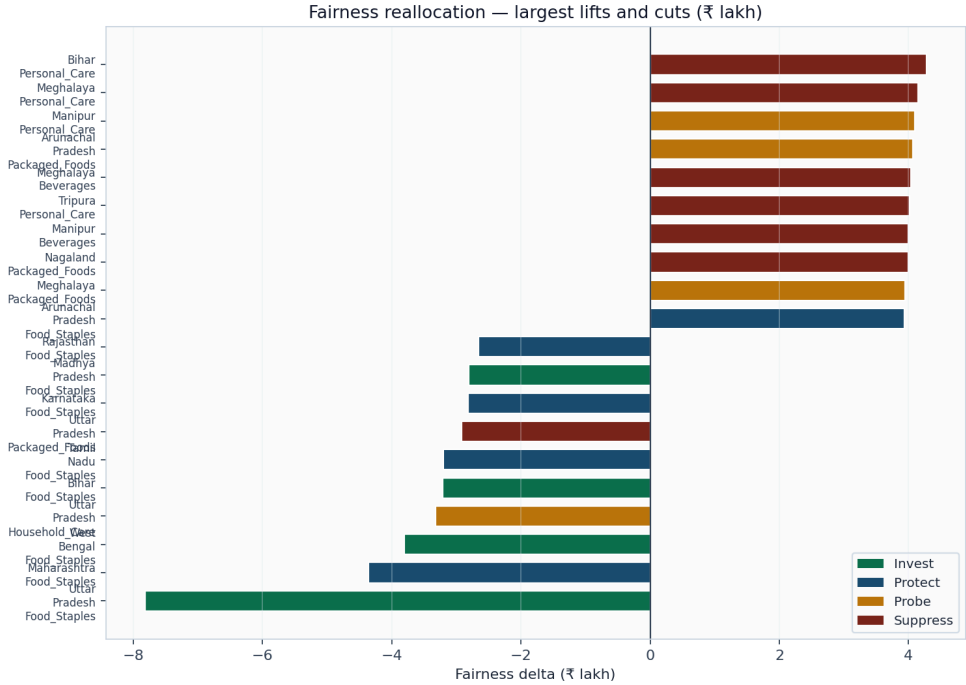


Fig. 6.6. Largest fairness lifts and cuts

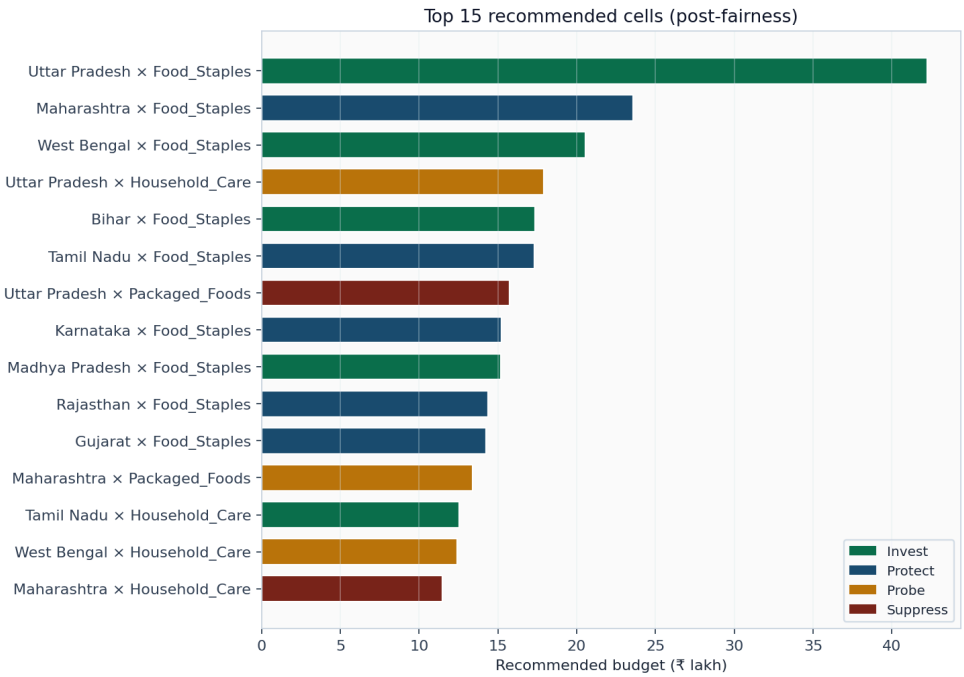


Fig. 6.7. Top recommended state × category cells

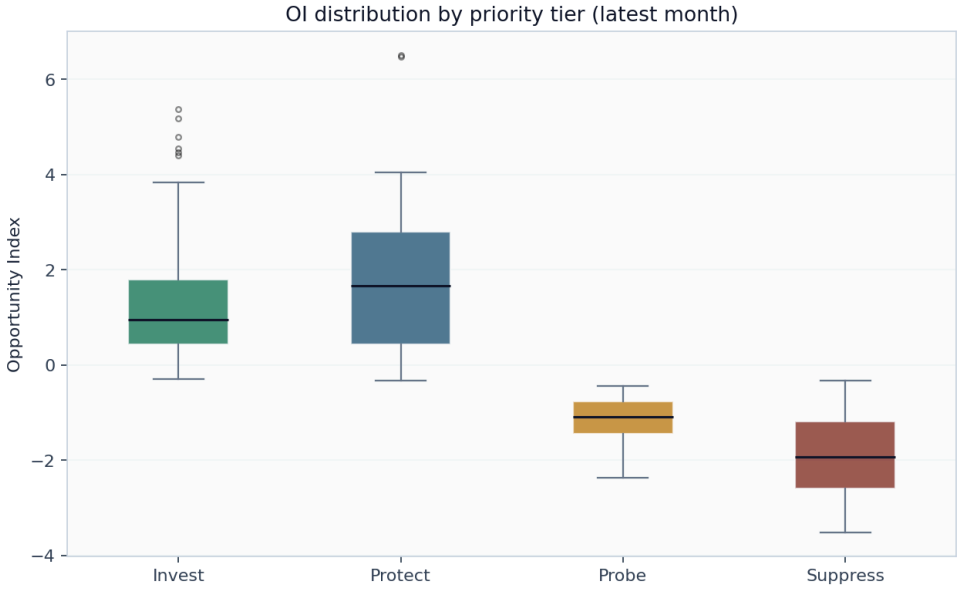


Fig. 6.8. OI distribution by priority tier

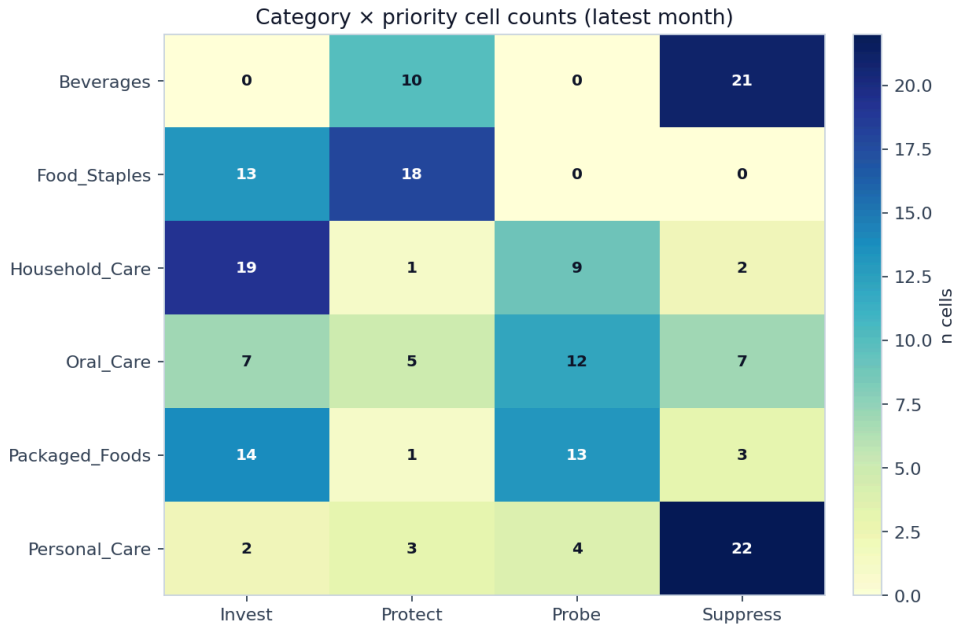


Fig. 6.9. Category × priority cell counts

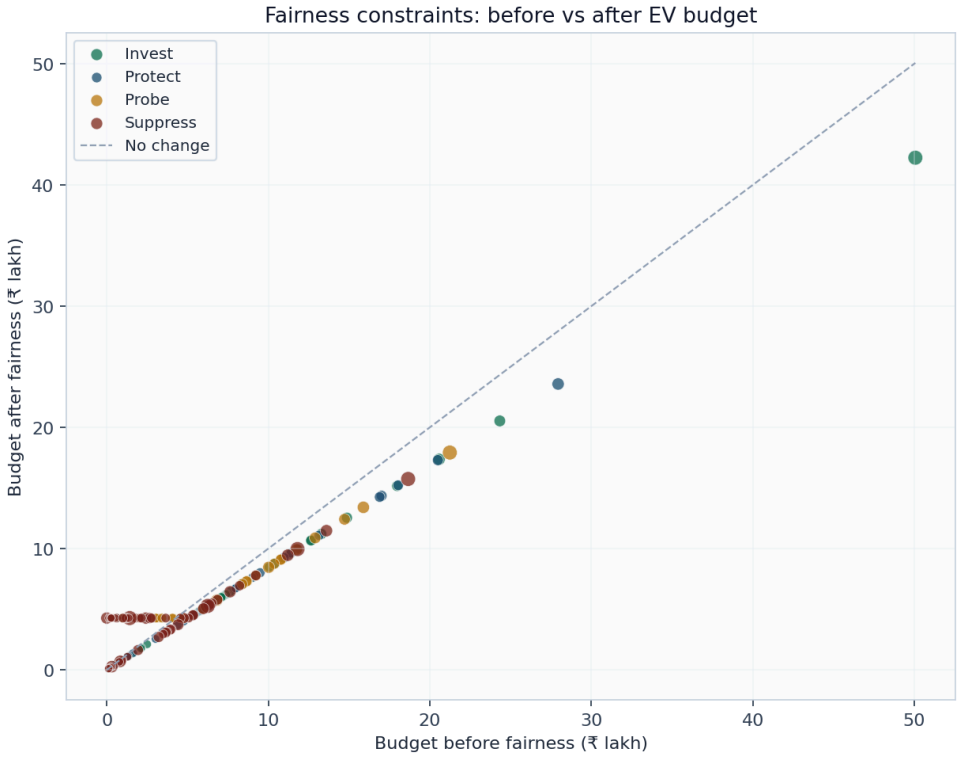


Fig. 6.10. Budget before vs after fairness

Live analytics tables

Analytics tables generated on 2026-07-31 from the Decision Stack run.

These tables accompany Results and findings. Narrative interpretation is in that chapter; numbers here are regenerated from the latest pipeline artefacts. This bind uses a harmonized secondary panel assembled from published official statistics (Appendix-I).

Panel reliability (missingness & cardinality)

Table 6.13. Panel reliability (missingness & cardinality).

column	missing_pct	n_unique
state	0	31.000
state_code	0	31.000
category	0	6.000
year_month	0	48.000
cpi_index	0	4,991
cpi_rural	0	4,960
cpi_urban	0	4,999
cpi_ru_gap	0	438.000
category_intensity	0	168.000
sales_share_proxy	0	8,644
population_million	0	31.000
urban_share	0	31.000
gsdp_pc_inr_thousand	0	112.000
broadband_per_100	0	1,232
wireless_per_100	0	886.000
rural_wage_inr	0	1,086
hces_mpce_share	0	164.000
festival_flag	0	2.000
festival_intensity	0	3.000
ad_clutter_index	0	6.000
expected_demand	0	186.000
unmet_demand	0	8,886
affordability	0	1,430
channel_readiness	0	1,358

OI × propensity priority matrix (latest-month cells)

Table 6.14. OI × propensity priority matrix (latest-month cells).

PriorityCell	n_cells	mean_OI	mean_propensity	mean_population
Invest	55	1.39	0.544	26.995
Probe	38	-1.143	0.541	44.305
Protect	38	1.855	0.51	25.768
Suppress	55	-1.882	0.514	56.449

Budget by PriorityCell

Table 6.15. Budget by PriorityCell.

PriorityCell	n_cells	recommended_budget	mean_budget	mean_OI	mean_propensity	budget_share_pct
Invest	55	32,252,622.63	586,411.32	1.39	0.544	32.253
Protect	38	21,676,596.82	570,436.76	1.855	0.51	21.677
Probe	38	22,616,644.82	595,174.86	-1.143	0.541	22.617
Suppress	55	23,454,135.73	426,438.83	-1.882	0.514	23.454

Fairness-constrained EV budget — top 40 cells

Table 6.16. Fairness-constrained EV budget — top 40 cells.

state	category	PriorityCell	OI	growth_propensity	recommended_budget	fairness_delta
Uttar Pradesh	Food_Staples	Invest	1.274	0.537	4,224,756.91	-781,781.51
Maharashtra	Food_Staples	Protect	1.534	0.505	2,358,279.37	-436,394.15
West Bengal	Food_Staples	Invest	1.545	0.54	2,053,622.92	-380,018.18
Uttar Pradesh	Household_Care	Probe	-1.512	0.543	1,792,572.83	-331,711.46

state	category	PriorityCell	OI	growth_propensity	recommended_budget	fairness_delta
Bihar	Food_Staples	Invest	0.256	0.537	1,736,832.96	-321,396.93
Tamil Nadu	Food_Staples	Protect	2.208	0.51	1,730,090.59	-320,149.27
Uttar Pradesh	Packaged_Foods	Suppress	-1.716	0.531	1,574,806.48	-291,414.30
Karnataka	Food_Staples	Protect	2.528	0.501	1,522,846.98	-281,799.32
Madhya Pradesh	Food_Staples	Invest	1.166	0.542	1,517,278.56	-280,768.89
Rajasthan	Food_Staples	Protect	1.283	0.531	1,436,215.46	-265,768.35

state	category	PriorityCell	OI	growth_propensity	recommended_budget	fairness_delta
Gujarat	Food_Staples	Protect	2.162	0.505	1,425,864.29	-263,852.89
Maharashtra	Packaged_Foods	Probe	-0.811	0.535	1,340,903.34	-248,131.07
Tamil Nadu	Household_Care	Invest	0.432	0.536	1,255,112.04	-232,255.59
West Bengal	Household_Care	Probe	-0.506	0.549	1,242,373.65	-229,898.38
Maharashtra	Household_Care	Suppress	-1.18	0.531	1,148,070.69	-212,447.83
Tamil Nadu	Packaged_Foods	Invest	-0.019	0.54	1,121,332.84	-207,500.05

state	category	PriorityCell	OI	growth_propensity	recommended_budget	fairness_delta
Andhra Pradesh	Food_Staples	Protect	1.807	0.512	1,109,204.16	-205,255.67
West Bengal	Packaged_Foods	Probe	-0.824	0.538	1,088,764.31	-201,473.32
Karnataka	Household_Care	Invest	0.438	0.538	1,068,923.53	-197,801.83
Karnataka	Packaged_Foods	Invest	0.398	0.542	1,066,824.32	-197,413.38
Kerala	Food_Staples	Protect	3.752	0.501	999,654.60	-184,983.77
Uttar Pradesh	Oral_Care	Suppress	-2.362	0.523	996,677.15	-184,432.80

state	category	PriorityCell	OI	growth_propensity	recommended_budget	fairness_delta
West Bengal	Oral_Care	Probe	-1.07	0.54	992,700.37	-183,696.91
Telangana	Food_Staples	Protect	3.036	0.51	961,925.09	-178,002.02
Gujarat	Household_Care	Invest	0.095	0.534	958,669.56	-177,399.59
Maharashtra	Oral_Care	Suppress	-1.576	0.526	945,954.57	-175,046.71
Karnataka	Oral_Care	Protect	-0.04	0.529	925,372.38	-171,238.02
Tamil Nadu	Oral_Care	Probe	-0.652	0.536	910,920.74	-168,563.78

state	category	PriorityCell	OI	growth_propensity	recommended_budget	fairness_delta
Madhya Pradesh	Household_Care	Probe	-0.749	0.55	909,310.40	-168,265.79
Gujarat	Packaged_Foods	Invest	-0.261	0.543	879,157.40	-162,686.05
Rajasthan	Household_Care	Probe	-0.622	0.536	874,679.04	-161,857.34
Kerala	Household_Care	Invest	2.222	0.538	847,463.29	-156,821.13
Bihar	Household_Care	Probe	-1.711	0.546	846,304.29	-156,606.65
Odisha	Food_Staples	Protect	0.842	0.532	801,374.22	-148,292.45

state	category	PriorityCell	OI	growth_propensity	recommended_budget	fairness_delta
Andhra Pradesh	Household_Care	Invest	0.009	0.543	780,323.00	-144,396.97
Gujarat	Oral_Care	Suppress	-0.571	0.532	778,983.09	-144,149.02
Rajasthan	Packaged_Foods	Probe	-0.951	0.538	778,837.07	-144,122.00
Kerala	Packaged_Foods	Invest	1.757	0.535	775,959.63	-143,589.54
Kerala	Oral_Care	Protect	1.821	0.526	770,791.52	-142,633.19

state	category	PriorityCell	OI	growth_propensity	recommended_budget	fairness_delta
Telangana	Household_Care	Invest	1.374	0.543	765,365.29	-141,629.08

Lowest 15 budget cells (Suppress / Probe tail)

Table 6.17. Lowest 15 budget cells (Suppress / Probe tail).

state	category	PriorityCell	OI	growth_propensity	recommended_budget	fairness_delta
Arunachal Pradesh	Personal_Care	Suppress	-1.955	0.528	9,526.38	-1,762.83
Sikkim	Personal_Care	Protect	0.586	0.519	10,517.23	-1,946.19
Sikkim	Beverages	Protect	1.376	0.483	11,674.25	-2,160.29
Arunachal Pradesh	Beverages	Suppress	-1.169	0.467	12,644.55	-2,339.85

state	category	PriorityCell	OI	growth_propensity	recommended_budget	fairness_delta
Nagaland	Personal_Care	Probe	-1.962	0.535	13,716.65	-2,538.23
Sikkim	Packaged_Foods	Invest	1.834	0.536	14,162.89	-2,620.81
Mizoram	Personal_Care	Probe	-0.553	0.539	14,452.17	-2,674.34
Sikkim	Oral_Care	Protect	2.099	0.525	14,556.56	-2,693.66
Sikkim	Household_Care	Protect	2.659	0.53	16,153.80	-2,989.22
Sikkim	Food_Staples	Protect	3.797	0.472	17,024.10	-3,150.27

state	category	PriorityCell	OI	growth_propensity	recommended_budget	fairness_delta
Mizoram	Beverages	Protect	0.231	0.518	17,580.66	-3,253.26
Arunachal Pradesh	Oral_Care	Probe	-0.437	0.538	19,100.90	-3,534.58
Nagaland	Beverages	Suppress	-1.176	0.508	19,586.63	-3,624.46
Mizoram	Packaged_Foods	Invest	0.686	0.555	21,102.91	-3,905.04
Puducherry	Personal_Care	Invest	0.618	0.55	22,472.39	-4,158.46

Sensitivity summary

Table 6.18. Sensitivity summary.

scenario	n_cells	invest_budget	fairness_pool_delta	floor_pool_inr	n_fairness_eligible
baseline	186	32,252,622.63	14,392,058.93	40,000,000	79
uniform_oi_plus_0.5	186	32,252,622.63	14,392,058.93	40,000,000	79

scenario	n_cells	invest_budget	fairness_pool_delta	floor_pool_inr	n_fairness_eligible
food_staples_oi_plus_0.5	186	32,486,222.64	14,483,894.06	40,000,000	79
oi_affordability_heavy_weights	186	28,892,211.35	15,015,606.84	40,000,000	79
fairness_floor_pool_20pct	186	33,801,443.50	5,912,476.88	20,000,000	79
fairness_floor_pool_40pct	186	32,252,622.63	14,392,058.93	40,000,000	79
fairness_floor_pool_60pct	186	30,638,926.70	22,016,360.47	60,000,000	79
digital_cap_15pct	186	32,252,622.63	14,392,058.93	40,000,000	79
digital_cap_25pct	186	32,252,622.63	14,392,058.93	40,000,000	79
digital_cap_40pct	186	32,252,622.63	14,392,058.93	40,000,000	79

Fairness constraint metadata

Table 6.19. Fairness constraint metadata.

total_budget	n_cells	n_fairness_eligible	floor_per_eligible_inr	floor_pool_inr	gsdp_cutoff	broadband_cutof
100,000,000	186	79	506,329.11	40,000,000	135.7	32.9

Model metrics snapshot

Table 6.20. Selected model and run metadata (from `model_metrics.json`).

field	value
model_label	Logistic regression
roc_auc	0.761
avg_precision	0.752
n_train	8,184
n_holdout	372
data_mode	harmonized_secondary_panel
n_panel_rows	8,928
n_states	31
n_categories	6
holdout_months	3
total_budget_inr	100,000,000
response_lift	0.040
fairness.n_cells	186
fairness.n_fairness_eligible	79
fairness.max_digital_share_low_readiness	0.250
fairness.sum_after	100,000,000

Market playbooks

The following playbooks summarise Invest / Protect / Probe / Suppress postures and recommended ₹ shares from the Decision Stack run reported in Results and findings.

Market playbooks — index

Generated from the Decision Stack budget allocation (OI × propensity × fairness EV).

Tier	Budget share
Invest	32.3%
Protect	21.7%
Probe	22.6%
Suppress	23.5%

Playbook: 01 invest

Market playbook: Invest

Definition

Priority cell **Invest** from the Opportunity Index × growth propensity matrix (high/low split at latest-month medians). Unit of analysis: Indian **state × FMCG category**.

Metric	Value
Cells (n)	55
Mean OI	1.390
Mean growth propensity	0.544
Total recommended budget (₹)	32,252,623
Share of national budget	32.3%

Recommended action

Invest

Campaign brief

Scale brand and modern-trade coverage where OI and propensity both clear the median

Planning emphasis

- Fund at or above the EV ₹ share while GT/MT capacity is still catching demand

Avoid

- Freeze spend in seed cells that already clear both medians

Top cells in this tier (by post-fairness budget)

state	category	OI	growth_propensity	recommended_budget	fairness_delta
Uttar Pradesh	Food_Staples	1.274	0.537	4,224,757	-781,782
West Bengal	Food_Staples	1.545	0.54	2,053,623	-380,018
Bihar	Food_Staples	0.256	0.537	1,736,833	-321,397
Madhya Pradesh	Food_Staples	1.166	0.542	1,517,279	-280,769

state	category	OI	growth_propensity	recommended_budget	fairness_delta
Tamil Nadu	Household_Care	0.432	0.536	1,255,112	-232,256
Tamil Nadu	Packaged_Foods	-0.019	0.54	1,121,333	-207,500
Karnataka	Household_Care	0.438	0.538	1,068,924	-197,802
Karnataka	Packaged_Foods	0.398	0.542	1,066,824	-197,413

Assumptions

- Response lift used in EV: **4.0%**
- Illustrative national budget: ₹**100,000,000**
- Fairness floors and digital-share caps applied as stated policy parameters (Methodology; sensitivity in Results and findings)

Charts

Figures in this report include:

- Opportunity Index versus growth propensity (Results and findings)
- Fairness lifts and cuts (Results and findings)
- Budget before versus after fairness (Results and findings)

Market playbook generated from the Decision Stack run reported in Results and findings.

Playbook: 02 protect

Market playbook: Protect

Definition

Priority cell **Protect** from the Opportunity Index × growth propensity matrix (high/low split at latest-month medians). Unit of analysis: Indian **state × FMCG category**.

Metric	Value
Cells (n)	38
Mean OI	1.855
Mean growth propensity	0.510
Total recommended budget (₹)	21,676,597
Share of national budget	21.7%

Recommended action

Protect

Campaign brief

Hold continuity media with a GT + MT mix while near-term growth cools

Planning emphasis

- Prioritise festival windows and clutter checks before adding weight

Avoid

- Cut continuity solely to polish a one-quarter ROI slide

Top cells in this tier (by post-fairness budget)

state	category	OI	growth_propensity	recommended_budget	fairness_delta
Maharashtra	Food_Staples	1.534	0.505	2,358,279	-436,394
Tamil Nadu	Food_Staples	2.208	0.51	1,730,091	-320,149
Karnataka	Food_Staples	2.528	0.501	1,522,847	-281,799
Rajasthan	Food_Staples	1.283	0.531	1,436,215	-265,768
Gujarat	Food_Staples	2.162	0.505	1,425,864	-263,853
Andhra Pradesh	Food_Staples	1.807	0.512	1,109,204	-205,256
Kerala	Food_Staples	3.752	0.501	999,655	-184,984
Telangana	Food_Staples	3.036	0.51	961,925	-178,002

Assumptions

- Response lift used in EV: **4.0%**
- Illustrative national budget: ₹**100,000,000**
- Fairness floors and digital-share caps applied as stated policy parameters (Methodology; sensitivity in Results and findings)

Charts

Figures in this report include:

- Opportunity Index versus growth propensity (Results and findings)
- Fairness lifts and cuts (Results and findings)
- Budget before versus after fairness (Results and findings)

Market playbook generated from the Decision Stack run reported in Results and findings.

Playbook: 03 probe

Market playbook: Probe

Definition

Priority cell **Probe** from the Opportunity Index × growth propensity matrix (high/low split at latest-month medians). Unit of analysis: Indian **state × FMCG category**.

Metric	Value
Cells (n)	38
Mean OI	-1.143

Metric	Value
Mean growth propensity	0.541
Total recommended budget (₹)	22,616,645
Share of national budget	22.6%

Recommended action

Probe

Campaign brief

Limit spend to EV-tied pilots; use digital only where TRAI readiness allows

Planning emphasis

- Read CPI and volume proxies from the pilot before any national scale-up

Avoid

- Treat a single propensity score as a green light for national rollout

Top cells in this tier (by post-fairness budget)

state	category	OI	growth_propensity	recommended_budget	fairness_delta
Uttar Pradesh	Household_Care	-1.512	0.543	1,792,573	-331,711
Maharashtra	Packaged_Foods	-0.811	0.535	1,340,903	-248,131
West Bengal	Household_Care	-0.506	0.549	1,242,374	-229,898
West Bengal	Packaged_Foods	-0.824	0.538	1,088,764	-201,473
West Bengal	Oral_Care	-1.07	0.54	992,700	-183,697
Tamil Nadu	Oral_Care	-0.652	0.536	910,921	-168,564
Madhya Pradesh	Household_Care	-0.749	0.55	909,310	-168,266
Rajasthan	Household_Care	-0.622	0.536	874,679	-161,857

Assumptions

- Response lift used in EV: **4.0%**
- Illustrative national budget: **₹100,000,000**
- Fairness floors and digital-share caps applied as stated policy parameters (Methodology; sensitivity in Results and findings)

Charts

Figures in this report include:

- Opportunity Index versus growth propensity (Results and findings)

- Fairness lifts and cuts (Results and findings)
- Budget before versus after fairness (Results and findings)

Market playbook generated from the Decision Stack run reported in Results and findings.

Playbook: 04 suppress

Market playbook: Suppress

Definition

Priority cell **Suppress** from the Opportunity Index × growth propensity matrix (high/low split at latest-month medians). Unit of analysis: Indian **state × FMCG category**.

Metric	Value
Cells (n)	55
Mean OI	-1.882
Mean growth propensity	0.514
Total recommended budget (₹)	23,454,136
Share of national budget	23.5%

Recommended action

Suppress

Campaign brief

Maintain thin presence or pause paid; keep the fairness floor only when policy requires it

Planning emphasis

- Reallocate residual ₹ into Invest / Protect once the floor is met

Avoid

- Load digital weight into cells that score low on both OI and propensity

Top cells in this tier (by post-fairness budget)

state	category	OI	growth_propensity	recommended_budget	fairness_delta
Uttar Pradesh	Packaged_Foods	-1.716	0.531	1,574,806	-291,414
Maharashtra	Household_Care	-1.18	0.531	1,148,071	-212,448
Uttar Pradesh	Oral_Care	-2.362	0.523	996,677	-184,433
Maharashtra	Oral_Care	-1.576	0.526	945,955	-175,047
Gujarat	Oral_Care	-0.571	0.532	778,983	-144,149
Andhra Pradesh	Oral_Care	-0.32	0.533	694,089	-128,440

state	category	OI	growth_propensity	recommended_budget	fairness_delta
Bihar	Packaged_Foods	-2.091	0.526	644,469	-119,257
Gujarat	Beverages	-1.145	0.491	578,784	-107,103

Assumptions

- Response lift used in EV: **4.0%**
- Illustrative national budget: ₹**100,000,000**
- Fairness floors and digital-share caps applied as stated policy parameters (Methodology; sensitivity in Results and findings)

Charts

Figures in this report include:

- Opportunity Index versus growth propensity (Results and findings)
- Fairness lifts and cuts (Results and findings)
- Budget before versus after fairness (Results and findings)

Market playbook generated from the Decision Stack run reported in Results and findings.

Playbook: cell 01 uttar pradesh food staples

Market playbook: Uttar Pradesh × Food_Staples

Field	Value
Priority cell	**Invest**
OI	1.274
Growth propensity	0.537
Population (mn)	199.8
GSDP pc (₹ '000)	82.6
Broadband / 100	29.1
Budget before fairness (₹)	5,006,538
Budget after fairness (₹)	4,224,757
Fairness delta (₹)	-781,782
Assumed digital share	25%
National priority rank	1

Action

Invest — Scale brand and modern-trade coverage where OI and propensity both clear the median

Planning emphasis

- Fund at or above the EV ₹ share while GT/MT capacity is still catching demand

Avoid

- Freeze spend in seed cells that already clear both medians

Cell dossier from the Decision Stack on the harmonized secondary panel.

Playbook: cell 02 maharashtra food staples

Market playbook: Maharashtra × Food_Staples

Field	Value
Priority cell	**Protect**
OI	1.534
Growth propensity	0.505
Population (mn)	112.4
GSDP pc (₹ '000)	265.5
Broadband / 100	57.6
Budget before fairness (₹)	2,794,674
Budget after fairness (₹)	2,358,279
Fairness delta (₹)	-436,394
Assumed digital share	30%
National priority rank	2

Action

Protect — Hold continuity media with a GT + MT mix while near-term growth cools

Planning emphasis

- Prioritise festival windows and clutter checks before adding weight

Avoid

- Cut continuity solely to polish a one-quarter ROI slide

Cell dossier from the Decision Stack on the harmonized secondary panel.

Playbook: cell 02 uttar pradesh household care

Market playbook: Uttar Pradesh × Household_Care

Field	Value
Priority cell	**Invest**
OI	1.648
Growth propensity	0.581
Population (mn)	199.8
GSDP pc (₹ '000)	78.4
Broadband / 100	27.9

Field	Value
Budget before fairness (₹)	4,606,999
Budget after fairness (₹)	3,833,720
Fairness delta (₹)	-773,279
Assumed digital share	25%
National priority rank	2

Action

Invest — Scale above-the-line + modern trade / e-com in high-OI high-growth cells

Do

- Fund brand-building and distribution where affordability and readiness align

Don't

- Under-invest relative to EV share; starve seed markets prematurely

Auto-generated from the Decision Stack (harmonized secondary panel).

Playbook: cell 02 uttar pradesh packaged foods

Market playbook: Uttar Pradesh × Packaged_Foods

Field	Value
Priority cell	**Probe**
OI	-1.284
Growth propensity	0.526
Population (mn)	199.8
GSDP pc (₹ '000)	78.4
Broadband / 100	27.9
Budget before fairness (₹)	2,433,705
Budget after fairness (₹)	2,073,788
Fairness delta (₹)	-359,916
Assumed digital share	25%
National priority rank	2

Action

Probe — Test-and-learn budgets; digital-first where readiness allows

Do

- Small EV-tied pilots; measure CPI/volume response proxies

Don't

- National rollout before propensity evidence

Auto-generated from the Decision Stack (harmonized secondary panel).

Playbook: cell 03 maharashtra food staples

Market playbook: Maharashtra × Food_Staples

Field	Value
Priority cell	**Protect**
OI	1.253
Growth propensity	0.487
Population (mn)	112.4
GSDP pc (₹ '000)	252.0
Broadband / 100	57.3
Budget before fairness (₹)	2,418,233
Budget after fairness (₹)	2,060,605
Fairness delta (₹)	-357,628
Assumed digital share	30%
National priority rank	3

Action

Protect — Defend share with continuity media; fortify GT + MT mix

Do

- Maintain presence; win festival windows; monitor clutter

Don't

- Cut continuously for short-term ROI optics

Auto-generated from the Decision Stack (harmonized secondary panel).

Playbook: cell 03 uttar pradesh packaged foods

Market playbook: Uttar Pradesh × Packaged_Foods

Field	Value
Priority cell	**Probe**
OI	-1.373
Growth propensity	0.560
Population (mn)	199.8

Field	Value
GSDP pc (₹ '000)	78.4
Broadband / 100	29.5
Budget before fairness (₹)	2,385,719
Budget after fairness (₹)	2,037,520
Fairness delta (₹)	-348,199
Assumed digital share	25%
National priority rank	3

Action

Probe — Test-and-learn budgets; digital-first where readiness allows

Do

- Small EV-tied pilots; measure CPI/volume response proxies

Don't

- National rollout before propensity evidence

Auto-generated from the Decision Stack (harmonized secondary panel).

Playbook: cell 03 west bengal food staples

Market playbook: West Bengal × Food_Staples

Field	Value
Priority cell	**Invest**
OI	1.545
Growth propensity	0.540
Population (mn)	91.3
GSDP pc (₹ '000)	141.6
Broadband / 100	38.5
Budget before fairness (₹)	2,433,641
Budget after fairness (₹)	2,053,623
Fairness delta (₹)	-380,018
Assumed digital share	45%
National priority rank	3

Action

Invest — Scale brand and modern-trade coverage where OI and propensity both clear the median

Planning emphasis

- Fund at or above the EV ₹ share while GT/MT capacity is still catching demand

Avoid

- Freeze spend in seed cells that already clear both medians

Cell dossier from the Decision Stack on the harmonized secondary panel.

Playbook: cell 04 uttar pradesh household care

Market playbook: Uttar Pradesh × Household_Care

Field	Value
Priority cell	**Probe**
OI	-1.512
Growth propensity	0.543
Population (mn)	199.8
GSDP pc (₹ '000)	82.6
Broadband / 100	29.1
Budget before fairness (₹)	2,124,284
Budget after fairness (₹)	1,792,573
Fairness delta (₹)	-331,711
Assumed digital share	25%
National priority rank	4

Action

Probe — Limit spend to EV-tied pilots; use digital only where TRAI readiness allows

Planning emphasis

- Read CPI and volume proxies from the pilot before any national scale-up

Avoid

- Treat a single propensity score as a green light for national rollout

Cell dossier from the Decision Stack on the harmonized secondary panel.

Playbook: cell 04 uttar pradesh oral care

Market playbook: Uttar Pradesh × Oral_Care

Field	Value
Priority cell	**Invest**
OI	0.968

Field	Value
Growth propensity	0.511
Population (mn)	199.8
GSDP pc (₹ '000)	78.4
Broadband / 100	27.9
Budget before fairness (₹)	3,502,462
Budget after fairness (₹)	2,914,578
Fairness delta (₹)	-587,884
Assumed digital share	25%
National priority rank	4

Action

Invest — Scale above-the-line + modern trade / e-com in high-OI high-growth cells

Do

- Fund brand-building and distribution where affordability and readiness align

Don't

- Under-invest relative to EV share; starve seed markets prematurely

Auto-generated from the Decision Stack (harmonized secondary panel).

Playbook: cell 04 west bengal food staples

Market playbook: West Bengal × Food_Staples

Field	Value
Priority cell	**Invest**
OI	1.792
Growth propensity	0.551
Population (mn)	91.3
GSDP pc (₹ '000)	134.4
Broadband / 100	39.6
Budget before fairness (₹)	2,309,123
Budget after fairness (₹)	1,972,103
Fairness delta (₹)	-337,020
Assumed digital share	45%
National priority rank	4

Action

Invest — Scale above-the-line + modern trade / e-com in high-OI high-growth cells

Do

- Fund brand-building and distribution where affordability and readiness align

Don't

- Under-invest relative to EV share; starve seed markets prematurely

Auto-generated from the Decision Stack (harmonized secondary panel).

Playbook: cell 05 bihar food staples

Market playbook: Bihar × Food_Staples

Field	Value
Priority cell	**Invest**
OI	0.256
Growth propensity	0.537
Population (mn)	104.1
GSDP pc (₹ '000)	59.0
Broadband / 100	24.9
Budget before fairness (₹)	2,058,230
Budget after fairness (₹)	1,736,833
Fairness delta (₹)	-321,397
Assumed digital share	25%
National priority rank	5

Action

Invest — Scale brand and modern-trade coverage where OI and propensity both clear the median

Planning emphasis

- Fund at or above the EV ₹ share while GT/MT capacity is still catching demand

Avoid

- Freeze spend in seed cells that already clear both medians

Cell dossier from the Decision Stack on the harmonized secondary panel.

Playbook: cell 05 uttar pradesh beverages

Market playbook: Uttar Pradesh × Beverages

Field	Value
Priority cell	**Invest**
OI	0.753
Growth propensity	0.508
Population (mn)	199.8
GSDP pc (₹ '000)	78.4
Broadband / 100	27.9
Budget before fairness (₹)	3,309,291
Budget after fairness (₹)	2,753,830
Fairness delta (₹)	-555,460
Assumed digital share	25%
National priority rank	5

Action

Invest — Scale above-the-line + modern trade / e-com in high-OI high-growth cells

Do

- Fund brand-building and distribution where affordability and readiness align

Don't

- Under-invest relative to EV share; starve seed markets prematurely

Auto-generated from the Decision Stack (harmonized secondary panel).

CHAPTER 7
MANAGERIAL IMPLICATIONS

7.1 Planning use

In national-level FMCG planning, we do not wait for the perfect vintage of data to emerge. Freshen the public panels driving the stack: CPI, TRAI readiness, GSDP, and rural wage. Also incorporate the HCES prior with every new wave. Repeat the classifier bakeoff before releasing any ₹ forecast.

AUC of the holdout and average precision are the quality checkpoint. On the reported run, discussion threshold will be AUC $\geq$ 0.60. Ranks are used when the chosen classifier passes the threshold.

Prioritize the key cells for the coming month in planning. Distribute expected value along with fairness-based allocations to the brand, category, and sales leads. Media, trade, and distribution then adopt the Invest / Protect / Probe / Suppress approach from the game books. This works well for quarterlies as well as for the pre-festival window.

7.1.1 Operational deployment path

The stack operates as a periodic planning process in a corporate environment where public extract data resides in either a warehouse or lakehouse table. The pipeline updates the state × category × month pane. The bake-off and allocator algorithms create priority and ₹ tables. Subsequently, a BI layer produces cell ranking and playbook content for brand and sales owners. This could be Power BI or any other dashboard, such as SAP Analytics or Databricks SQL. Governance remains within category or board level. Fairness percentiles, floor-pool allocation, digital cap, and lift l continue being policy attributes that planners can edit prior to the subsequent stack run. ERP or SAP systems of media commitments consume the approved allocations. The research contribution ends at the auditable recommendation surface, not at the media booking automation level.

7.2 Priority-cell stance

These four priorities are all stances, not predictions.

Invest markets are those where the intersection of growth and growth propensity occurs. These markets can be scaled on brand and distribution basis. Go for above-the-line and modern trade/e-com to the extent that channel readiness enables that. Maintain funding at or above the level of EV to prevent any early pruning of that signal.

In contrast, protection is required in high OI markets whose growth momentum is waning. Such markets require continuity media, better general and modern trade as well as festivals that are higher than short-term ROI optics.

Probing markets require proofing before scaling. Conduct small pilots, conduct digital only where channel readiness enables that, and scale only once the pilot is clean.

In suppress markets, reduce spending wherever possible, maintain fairness floors wherever such is dictated by policy and shift remaining budget to invest and protect markets.

Table 7.1. Managerial stance by priority cell.

Priority cell	Management stance	Media / channel posture	Governance
Invest	Scale brand + distribution where OI and growth propensity align	Above-the-line + modern trade / e-com where readiness allows	Fund at or above EV share. Avoid premature cuts

Priority cell	Management stance	Media / channel posture	Governance
Protect	Hold high-OI cells whose growth has cooled	Continuity media. Fortify GT + MT. Win festival windows	Resist cutting for short-term ROI optics
Probe	Learn before scaling	Small EV-tied pilots. Digital where readiness allows	Confirm evidence before national rollout
Suppress	Keep a thin presence	Pause paid where feasible. Retain fairness floor when policy requires it	Move residual budget to Invest/Protect

7.3 Current allocation

The bake-off winner is locked down even in Chapter 6. This is also true of the Invest allocation share of the illustrative budget of ₹10 crore. So too are fairness reallocation allocations on the latest-month surface. For planning purposes, however, the issue is what to make of that surface. Growth and geographical penetration come in dual layers. Expected-value ranks direct spending towards high-ranked cells. Fairness floors ensure a rule-driven presence in thin markets. They don't jettison the expected value structure. Where the TRAI readiness is poor, the digital cap (25% in the reported run) tilts execution towards more traditional trade and vernacular channels. Some of the digital spending remains, however.

Uttar Pradesh × Food Staples exemplifies the dual-layer principle. On the reported surface, it lies in Invest (OI ≈ 1.27, propensity ≈ 0.54). Pre-fairness EV puts it among the top national choices. With the floors included, the recommended share is around ₹42 lakh. The cell still provides for growth. Some of the unconstrained mass is utilized in satisfying the floors in other cells. Boards can use this argument cell-by-cell. They do not have to refer only to a single national average.

7.4 Market playbooks

Tier dossiers and top-cell dossiers turn the same surface into meeting-ready language. Each dossier answers four questions. Which cell is in view? What ₹ share does it carry? What campaign posture fits that stance? Which actions should pause? Analytics output then stops being a one-off slide. It becomes a routine input to planning meetings.

7.5 Budget figures

Budget shares in this chapter are illustrative decision support for national planning. They rest on the assumed response lift. They also rest on the panel documented earlier. Teams that need causal ROI, promotion elasticity, or creative tests can add experiments later. They can also add brand-data marketing-mix work. The stack itself stops at auditable ranking and allocation under stated assumptions.

CHAPTER 8
CONCLUSION, FUTURE SCOPE AND SOCIAL IMPACT

8.1 Conclusions

Geo-marketing effort finds geographies with potential markets. Expected-response budgeting budgets on the basis of expected response. Inclusive-growth practice calls for floors in thinner geographies. However, the corpus considered does not tie these threads into one single layer based on India state × category data with auditable ₹ shares at the end. The present study plugged this missing link through a Market Decision Stack on harmonized secondary data.

Market Decision Stack accomplished three tasks within one pass. First, it calculated Opportunity Index, independent of near-term growth propensity. Second, it chose the propensity model via time-holdout bake-off test; logistic regression was chosen (ROC-AUC for holdout data 0.761; average precision 0.752). Third, it transformed the layer into Invest/Protect/Probe/Suppress cells, and further into expected-value ₹ shares under fairness floors and digital share cap. In the illustrative latest-month ₹10 crore surface, Invest took up nearly one-third of the budget, and fairness shifted around ₹1.44 crore to around 79 thinner cells without disturbing EV order. All tables are available in Chapter 6.

The innovation in this case is thus applied packaging, like the distinction made by Gupta and Steenburgh between predict-then-allocate, by Tudor between opportunity mapping and spend rules, and in Celis et al.'s approach, where fairness is an explicit constraint rather than a latent residual. No new family of classifiers has been developed here. Media ROI in causal terms is out of the picture. Ditto a social optimum for fairness. The actionable point regarding Indian FMCG planning is limited to using public aggregates to trace a verifiable route from portal sums to prioritized cells, rupee allocations, and playbooks to argue about.

8.2 Limitations

The data proxies abound in abundance. Public CPI, GSDP, and TRAI series proxy opportunity and propensity respectively. Sell-out and media logs are not in this set. The public data can be sometimes out-of-sync with what retailers do. The planning month can depend on slightly outdated opportunity and/or propensity.

The current estimates combine with published magnitude inputs in the Decision Stack framework. New portal micro-data have to be pulled in and validated (Appendix-I). The state nicknames and CPI→FMCG category mapping are yet to be validated before a committee finalizes the taxonomy. The category mapping simplifies public CPI categories down to six planning categories.

Design decisions constrain interpretation too. Equal OI weights remain transparent but might overlook unique board considerations. The affordability-centric bias in Chapter 6 captures how ₹ shares can move even when indices remain highly correlated. Response lift stays fixed at (= 0.04) on this panel. Under the EV share formula, it does not reorder cells. Fairness floors, eligibility percentiles, and the digital cap are policy knobs: change them and the ₹ table moves. That is why varied stress tests stay in the method. Public aggregates support ranking and ruled allocation. They do not identify causal media effects.

Validation is internal. It includes time holdout, continuous-vs.-classifier cross-validation, calibration, walk-forward and 2024->2025 calendar splits, as well as stress tables. No industrial sell-out backtest is performed. No later-than-panel vintage is claimed. No district-level pilot test is performed. These belong to Tier-3 future work.

Delivery is report-first. The compiled report and playbooks ship with this study. An interactive UI is optional later work. Cell logic uses within-month z-scores and stated allocation rules. A full spatial-autoregressive dependence model waits on denser brand series if those appear. CPI lag-1, lag-3, and year-on-year terms are pre-specified horizons. ACF/PACF and nested AUC diagnostics in Appendix-I (Section I.7) back that choice. Festival and some digital coefficients stay weak in the linear logit even though the joint model ranks well. Table 6.4 shows one-way links on this panel, not a causal feature story.

Table 8.1. Key risks and mitigations applied in this study.

Risk	Mitigation in this study
Reading ₹ shares as causal media ROI	Assumed lift. Illustrative decision-support framing
Silent bad joins	Canonical state codes + reliability report
Reading fairness floors as a unique optimum	Sensitivity tables. Policy framing
Reading equal weights as the only possible OI	Weight-tilt correlation + budget sensitivity
Reading harmonized figures as official microdata cells	Dataset label. Appendix-I. Portal citations
Over-reading uniform sensitivity	Discriminative Food_Staples / floor-pool / weight checks reported
Treating holdout AUC as external validation	Explicit internal-check scope. Tier-2 cross-checks reported. Industrial / district pilots listed as future work

8.3 Future scope

A natural next step is to replace selected joins with byte-identical MoSPI, RBI, TRAI, and Census portal downloads, recording pull dates in Appendix-I. An anonymised brand-volume pilot for a few categories could then test whether EV ranks align with internal movement—industrial verification that this public panel cannot supply. A later vintage (train through 2025; score 2026 months once published) would extend the calendar holdout beyond the present freeze. Where open district series exist, a single-state pilot could check whether state ranks persist at finer grain. ARIMA and seasonal-naive baselines on a few state–category series would cross-check forecasts without displacing propensity ranking inside the Decision Stack.

8.4 Social impact

Social impact here is not a measured welfare gain for households. It is the practical consequence of putting inclusive geographic rules on the same surface as growth ranks. National FMCG planning can then keep a ruled presence outside already-thick metros while still funding high-opportunity cells.

8.4.1 Direct impact — inclusive geographic reach

On the reported ₹10 crore surface, 79 cells met the low-GSDP or low-penetration rule. A ₹4.0 crore floor pool (≈ ₹5.06 lakh per eligible cell) produced about ₹1.44 crore of positive fairness reallocation. Eligible cells held roughly half of post-fairness budget.

Expected-value logic concentrates spend where population, propensity, and positive opportunity align. That is commercially sound. It is also uneven. Thin state×category cells can drop out of the plan even when households there still buy staples and household care. Fairness floors enter as a stated policy lever, not as a unique social optimum. Uttar Pradesh × Food_Staples stayed near the top after floors (≈ ₹42 lakh). Part of unconstrained Invest mass funded thinner cells. Growth spend and lagging-market presence then sit as two layers on one table. Boards can defend both without inventing a separate “CSR budget” sheet.

Where TRAI broadband readiness fell below the cutoff, digital share was capped at 25%. That keeps some of the plan on traditional retail and vernacular media in low-readiness states, so households and kirana networks there are not reached only through digital channels they may not use reliably.

8.4.2 Indirect impact — transparency, skills, and planning culture

All Opportunities Index inputs and panel joins can be found in official public data (see Chapter 3 and Appendix-I). Examiners, regulators, and brand managers can contest a cell's rating without resorting to confidential sellouts. This accountability is a social benefit indirectly derived from the fact that vendor ratings would block the joins.

Market strategies translate Invest/Protect/Probe/Suppress into meeting language. Category managers, sales leaders, and regional teams who do not develop classifiers themselves will still be able to debate which thin markets should be held down and which suppress cells are thin deliberately. Planning culture evolves from national averages to a per-cell perspective, which resembles more the language of growth and aspirational regions used in India rather than the metros-only media buys.

The stack also lowers the bar for universities, MSME brands, and state agencies that lack CMIE-grade panels. A repeat path from CPI, GSDP, TRAI, wages, and HCES priors to ruled ₹ shares works for teaching, pilot planning, and NGO or public-sector category work where brand PII cannot be shared.

8.4.3 Impact summary

Table 8.2. Social impact pathways linked to Decision Stack outputs on the reported run.

Pathway	Who is affected	Mechanism in this study	Observable impact on the reported surface
Geographic inclusion	Households and trade in thinner / lower-GSDP state×category cells	Fairness floors (H11) on EV shares	≈ ₹1.44 crore positive reallocation. 79 eligible cells. ~half of post-fairness budget in eligible set
Channel access equity	Buyers and kirana networks where broadband readiness is weak	Digital-share cap (25%) under TRAI cutoff (H12)	Execution pushed toward GT / vernacular where readiness is thin
Essential-category presence	Staples and household-care demand outside large metros	Priority matrix + floors on Food_Staples and related cells	High-OI Invest cells retained. Floors keep lagging cells in the ₹ table
Decision transparency	Examiners, planners, civil society reviewers	Harmonized secondary panel + Appendix-I disclosure	Scores and joins can be audited without brand PII
Planning literacy	Brand, sales, and regional non-specialists	Auto-generated market playbooks	Stances discussable cell by cell in planning meetings
Open secondary-data access	Students, MSMEs, agencies without proprietary panels	Public-statistics Decision Stack	Repeat path from portals to ruled shares

Such impacts remain decision-support impacts under assumed response lift (= 0.04) and an illustrative national budget. Causal poverty reduction, consumer-surplus estimates, and employment multipliers of the fairness stage lie beyond the evidence on this panel. The tangible claim is narrower: a ruled, inspectable shift of planning attention and funds toward thinner markets, with a governance trail that keeps that shift discussable.

REFERENCES

Academic sources

1. Baray, J. (2025). Geomarketing and retail location. In Elgar Encyclopedia of Retailing. Edward Elgar. https://doi.org/10.4337/9781035319701.00207
2. Bawindsom, K. M., Rodrigue, K. W., Galiam, O. S. M., Paul, T., & Alfred, S. (2024). Enhancing broadband access in urban Burkina Faso: A GIS and machine learning-based geomarketing approach to FTTx potential and alternatives. In 2024 IEEE MNE3SD. https://doi.org/10.1109/MNE3SD63831.2024.10812152
3. Bradlow, E. T., Bronnenberg, B., Russell, G. J., Arora, N., Bell, D. R., Duvvuri, S. D., ter Hofstede, F., Sismeiro, C., Thomadsen, R., & Yang, S. (2005). Spatial models in marketing. Marketing Letters, 16(3/4), 267–278. https://doi.org/10.1007/s11002-005-0457-7
4. Bronnenberg, B. J., & Mahajan, V. (2001). Unobserved retailer behavior in multimarket data: Joint spatial dependence in market shares and promotion variables. Marketing Science, 20(3), 284–299.
5. Celis, L. E., Mehrotra, A., & Vishnoi, N. K. (2019). Toward controlling discrimination in online ad auctions. In Proceedings of the 36th International Conference on Machine Learning (PMLR 97). http://proceedings.mlr.press/v97/mehrotra19a/mehrotra19a.pdf
6. Cliquet, G. (2020). From geomarketing to spatial marketing. In Spatial Economics Volume II: Applications. Springer. https://doi.org/10.1007/978-3-030-40094-1_10
7. Cliquet, G. (2025). Retail location theories and models. In Elgar Encyclopedia of Retailing. Edward Elgar. https://doi.org/10.4337/9781035319701.00214
8. Cliquet, G., & Baray, J. (2020). Location-based Marketing: Geomarketing and Geolocation. Wiley / ISTE. https://doi.org/10.1002/9781119721338
9. Crisóstomo-Berrocal, T., Sierra-Liñan, F., & Carbonell-Michael, C. (2022). Digital platform based on geomarketing as an improvement in micro and small enterprises. Indonesian Journal of Electrical Engineering and Computer Science, 27(1), 395–403. https://doi.org/10.11591/ijeecs.v27.i1.pp395-403
10. De Chiara, D., Di Giovanni, P., Sebillo, M., Tortora, G., & Vitiello, G. (2011). Geomarketing policies and augmented reality for advertisement delivery on mobile devices. In Proceedings of DMS 2011. Knowledge Systems Institute.
11. Farris, P. W., Bendle, N. T., Pfeifer, P. E., & Reibstein, D. J. (2010). Marketing Metrics: The Definitive Guide to Measuring Marketing Performance (2nd ed.). Wharton School Publishing / Pearson.
12. Garcia, M. B. (2022). Location-based marketing using mobile geofencing: Lessons learned from a user-centred application development research. International Journal of Technology Marketing. https://doi.org/10.1504/ijtmkt.2023.127322
13. Gelauff, L., Goel, A., Munagala, K., & Yandamuri, S. (2020). Advertising for demographically fair outcomes. arXiv preprint arXiv:2006.03983. https://arxiv.org/abs/2006.03983
14. Gladilin, P., & Levina, P. (2020). Method for analyzing the location, assortment and success of outlets based on transactional data. In ACM International Conference Proceeding Series. https://doi.org/10.1145/3446999.3447005
15. Gupta, S., & Steenburgh, T. J. (2008). Allocating Marketing Resources (HBS Working Paper 08-069). Harvard Business School. https://www.hbs.edu/ris/Publication%20Files/08-069_17a7715d-c34b-4d9e-92fa-2ea2834a0cbe.pdf
16. Kerzina, E., Antineskul, E., & Patrusheva, L. (2023). Geomarketing in the placement of retail in the development of a million-plus city (on the example of a regional retail in Russia, Perm). Lecture Notes in Networks and Systems. https://doi.org/10.1007/978-3-031-28086-3_68
17. Kita, P., & Čvirik, M. (2024). Retailer processes centred on the food market as the main determinate of business models in the context of retail size. Journal of Retailing and Consumer Services, 79, 103937. https://doi.org/10.1016/j.jretconser.2024.103937

18. Kotler, P., Keller, K. L., Koshy, A., & Jha, M. (2013). Marketing Management: A South Asian Perspective (14th ed.). Pearson — geographic market prioritisation and Indian institutional context.
19. Lilien, G. L., Rangaswamy, A., & De Bruyn, A. (2013). Principles of Marketing Engineering (2nd ed.). DecisionPro — geographic and resource-allocation chapters.
20. NITI Aayog. (2018). Strategy for New India @ 75. Government of India — inclusive growth and aspirational / lagging-region framing that motivates fairness floors as a policy design input. https://www.niti.gov.in/
21. NITI Aayog. (2023). An Assessment of Aspirational Districts Programme. Government of India / Institute for Competitiveness — lagging-district uplift evaluation that motivates fairness-eligible markets as a policy design input. https://niti.gov.in/sites/default/files/2023-03/An-Assessment-Of-Aspirational-Districts-Programme.pdf
22. OECD / European Commission Joint Research Centre. (2008). Handbook on Constructing Composite Indicators: Methodology and User Guide. OECD Publishing. https://doi.org/10.1787/9789264043466-en
23. Opportunity Index. (2023). The 2021, 2022, & 2023 Opportunity Index Technical Supplement. https://opportunityindex.org/wp-content/uploads/2024/09/2023-Opportunity-Index-Technical-Supplement.pdf
24. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., … Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71. https://doi.org/10.1136/bmj.n71
25. Pandey, R. (2022). It's Happy Diwali for economy despite inflation, recession fears. But seasonal joy must sustain (moving-holiday / Diwali seasonality note). National Institute of Public Finance and Policy (NIPFP). https://nipfp.org.in/media/medialibrary/2022/10/21102022.pdf
26. Paul, J., & Criado, A. R. (2020). The art of writing literature review: What do we know and what do we need to know? International Business Review, 29(4), 101717. https://doi.org/10.1016/j.ibusrev.2020.101717
27. Poledníková, E. (2017). Weighting methods for constructing composite indices in regional development. In Proceedings of the International Scientific Conference Hradec Economic Days. https://msed.vse.cz/msed_2017/article/225-Polednikova-Eva-paper.pdf
28. Rossi, P. E., Allenby, G. M., & McCulloch, R. (2005). Bayesian Statistics and Marketing. Wiley — demand and hierarchical market models (conceptual backdrop).
29. Salgado, R. R., Libório, M. P., Da Silva Martin, C. A. P., & Bernardes, P. (2019). Two-stage geomarketing WebGIS solution for a retail location problem. RA'E GA — O Espaço Geográfico em Análise, 46(1). https://doi.org/10.5380/raega.v46i1.56067
30. Samak, P., Bhurke, V., Bacche, O., & Khan, A. (2026). Tradition meets market: Consumer behavior and brand competition in the Diwali sweets segment. International Journal for Multidisciplinary Research (IJFMR), 8(1). https://www.ijfmr.com/papers/2026/1/68345.pdf
31. Tamilselvi, R., & Rajeswari, K. (n.d.). A study on consumer perception and buying behaviour towards festive sale in India with special reference to Chennai city. The Online Journal of Distance Education and e-Learning (TOJDEL). https://www.tojned.net/journals/tojdel/articles/v11i02b/v11i02b-09.pdf
32. Tellis, G. J. (2004). Effective Advertising: Understanding When, How, and Why Advertising Works. Sage.
33. Tudor, C. (2025). A geospatial framework for retail suitability modelling and opportunity identification in Germany. ISPRS International Journal of Geo-Information, 14(9), 342. https://doi.org/10.3390/ijgi14090342
34. Tudor, C., Girlovan, A., & Botoroga, C.-A. (2025). Where business meets location intelligence: A bibliometric analysis of geomarketing research in retail. ISPRS International Journal of Geo-Information, 14(8), 282. https://doi.org/10.3390/ijgi14080282
35. Wichmann, J. R. K., Scholdra, T. P., & Reinartz, W. J. (2021). What drives inner city attractiveness for society? The role of brick-and-mortar stores. Review of Marketing Research, 18, 203–229. https://doi.org/10.1108/S1548-643520210000018011
36. Zhao, L. Q., van Duynhoven, A., & Dragićević, S. (2024). Machine learning for criteria weighting in GIS-based multi-criteria evaluation: A case study of urban suitability analysis.

Land, 13(8), 1288. https://doi.org/10.3390/land13081288

Public data portals

This study's working dataset is a harmonized secondary panel assembled from the published official statistics below.Retrieval date (core portal layers, Refs. 37–46): 2026-07-15. Retrieval date (safe-scope extensions, Refs. 47–48): 2026-07-26. Coverage us ed: 31 states/UTs × 6 FMCG categories × 48 months (8,928 panel rows; 2022-01–2025-12).Layer-to-variable mapping and citation examples: Appendix-I, Tables I.2–I.4.

1. Ministry of Statistics and Programme Implementation (MoSPI). (n.d.). Consumer Price Index (CPI) publications and download tables. Retrieved 2026-07-15 from https://www.mospi.gov.in/cpi
2. Ministry of Statistics and Programme Implementation (MoSPI). (n.d.). State Domestic Product / GSDP tables. Retrieved 2026-07-15 from https://www.mospi.gov.in/state-domestic-product
3. Ministry of Statistics and Programme Implementation (MoSPI). (n.d.). Download tables / data portal. Retrieved 2026-07-15 from https://www.mospi.gov.in/download-tables-data
4. Reserve Bank of India (RBI). (n.d.). Handbook of Statistics on Indian States. Retrieved 2026-07-15 from https://www.rbi.org.in/Scripts/AnnualPublications.aspx?head=Handbook%20of%20Statistics%20on%20Indian%20States
5. Telecom Regulatory Authority of India (TRAI). (n.d.). Performance Indicator Reports. Retrieved 2026-07-15 from https://www.trai.gov.in/release-publication/reports/performance-indicators-reports
6. Office of the Registrar General & Census Commissioner, India. (n.d.). Census of India portal. Retrieved 2026-07-15 from https://censusindia.gov.in/census.website/
7. Unique Identification Authority of India (UIDAI). (n.d.). Ecosystem / authentication statistics. Retrieved 2026-07-15 from https://uidai.gov.in/
8. International Institute for Population Sciences (IIPS). (n.d.). National Family Health Survey (NFHS). Retrieved 2026-07-15 from https://nfhsbiips.in/
9. National Portal of India. (n.d.). Calendar / public holiday references. Retrieved 2026-07-15 from https://www.india.gov.in/calendar
10. Advertising Standards Council of India (ASCI). (n.d.). Industry resources (category advertising norms / guidance). Retrieved 2026-07-15 from https://www.ascionline.in/
11. Labour Bureau, Ministry of Labour and Employment. (n.d.). Wage Rates in Rural India (WRRI). Retrieved 2026-07-26 from https://www.labourbureau.gov.in/wage-rates-in-rural-india
12. Ministry of Statistics and Programme Implementation (MoSPI). (n.d.). Household Consumption Expenditure Survey (HCES) (2022–23 / 2023–24 style releases; microdata / published tables). Retrieved 2026-07-26 from https://microdata.gov.in/

Citation note

Suggested in-text form when referring to this study's dataset:

This study analyses a harmonized secondary panel assembled from published official statistics (MoSPI CPI and GSDP; RBI Handbook of Statistics on Indian States; TRAI Performance Indicator Reports; Census of India / NFHS / UIDAI population summaries; National Portal festival calendar; ASCI-related ad-intensity aggregates; Labour Bureau WRRI rural wages; MoSPI HCES MPCE shares), retrieved 2026-07-15 (extensions 2026-07-26; Appendix-I; Refs. 37–48).

When byte-identical portal Excel dumps replace the harmonized secondary panel, update each portal entry with the new retrieval date and file title, e.g.:

Organisation (Year). Title. Retrieved from URL on YYYY-MM-DD.

Screen counts and the narrative-synthesis core are reported in the Literature review (Figure 2.1; Tables 2.1–2.2). Scopus exports are dated 2026-07-21; the Scholar/OA gap-fill is dated 2026-07-26; the methods OA fill (composite weighting; constrained fairness) is dated 2026-07-31.

APPENDIX-I
DATA SOURCES AND PANEL CONSTRUCTION

This appendix is the citable data list for this study. It records what data are used, who published them, where to find them online, when they were pulled, and how they were joined into the analysis panel.

I.1 Dataset designation

Table I.1. Dataset designation terms, definitions, and status in this bind.

Term	Definition	Status in this bind
Harmonized secondary panel	A longitudinal panel at state × category × month built by joining secondary data drawn from published official statistics (and closely related public aggregates) onto one analysis schema	Used for all reported results
Published official statistics	Public MoSPI, RBI, TRAI, Census of India, NFHS/UIDAI, Labour Bureau, and related releases cited in References 37–48	Source family for magnitudes
Primary survey / admin microdata	Field surveys or byte-identical portal Excel extracts collected/downloaded as individual records	Not used

Methods framing: this study is a secondary analysis of published totals, after geographic and category harmonisation.

Important: Cell-level values support repeat Opportunity Index / propensity / budget methods. They are not to be quoted as official government microdata cells.

I.2 How to cite the data

Use References 37–48 for the bibliography. Use this table when naming what entered the Decision Stack.

Table I.2. Analytical uses, variables, publishers, cites, URLs, and retrieval dates.

Analytical use in this study	Variables / constructs	Publisher	Bibliographic cite	URL
Inflation / category demand pressure	CPI-linked index, R/U gap, category intensity, sales-share proxy	MoSPI	Refs. 37, 39	https://www.mospi.gov.in/cpi · https://esankhyiki.mospi.gov.in/macroindicators?product=cp
Affordability / ability to pay	GSDP per capita (₹ thousand), rural wage blend	MoSPI, RBI, Labour Bureau	Refs. 38, 40	https://www.mospi.gov.in/state-domestic-product · https://www.rbi.org.in/Scripts/AnnualPublications.aspx?head=Handbook%20of%20Statistics%20on%20Indian%20S · https://www.labourbureau.gov.in/wage-rates-in-rural-india

Analytical use in this study	Variables / constructs	Publisher	Bibliographic cite	URL
Digital / channel readiness	Broadband per 100, wireless per 100	TRAI	Ref. 41	https://www.trai.gov.in/release-publication/reports/performa indicators-reports
Market size / urbanisation	Population (mn), urban share	Census of India, UIDAI, NFHS (IIPS)	Refs. 42, 43, 44	https://censusindia.gov.in/census.website/ · https://uidai.gov https://nfhsbiips.in/
Festival seasonality	Festival flag, festival intensity by month	National Portal of India	Ref. 45	https://www.india.gov.in/calendar
Competitive clutter	Category ad-clutter index	ASCI (+ published ad-ex aggregates)	Ref. 46	https://www.ascionline.in/
Rural wages (affordability IV)	Rural wage (₹/day) blend	Labour Bureau, RBI Handbook wage tables	Refs. 47, 40	https://www.labourbureau.gov.in/wage-rates-in-rural-india · https://www.rbi.org.in/Scripts/AnnualPublications.aspx?head=Handbook%20of%20Statistics%20on%20Indian%20S
HCES structural prior	Category MPCE share by state	MoSPI HCES	Ref. 48	https://microdata.gov.in/

In-text citation example: CPI-linked demand pressure uses MoSPI Consumer Price Index publications (Ref. 37). Affordability uses MoSPI GSDP tables and the RBI Handbook of Statistics on Indian States (Refs. 38, 40).

I.3 Panel coverage

Table I.3. Coverage dimensions of the harmonized secondary panel.

Dimension	Value used in reported results
Unit of analysis	State / UT × FMCG decision category × month
Geography	31 states and union territories
Categories	6 (Food_Staples, Beverages, Personal_Care, Household_Care, Packaged_Foods, Oral_Care)
Time span	48 months (2022-01–2025-12, excludes 2020–2021)
Panel rows	8,928 (= 31 × 6 × 48)
Latest-month decision cells	186 (= 31 × 6)
Dataset type	Harmonized secondary panel from published official statistics

I.4 Panel layers

Table I.4. Panel-layer inventory — encodings and source basis.

Panel layer	What it encodes	Source basis (cite)
CPI state × category × month	CPI-like index, R/U indices & gap, category intensity, sales-share proxy	MoSPI CPI / eSankhyiki (Refs. 37, 39) harmonized to Decision Stack schema
GSDP by state × year	GSDP per capita (₹ thousand)	MoSPI / RBI Handbook scale (Refs. 38, 40), joined to monthly panel
TRAI readiness by state × month	Broadband and wireless per 100	TRAI Performance Indicator Reports (Ref. 41)
Population by state	Population (mn) and urban share	Census / NFHS / UIDAI summaries (Refs. 42–44)
Festival calendar by month	Festival flag / intensity	National Portal calendar coding (Ref. 45)
Ad clutter by category	Ordinal clutter index	ASCI / published aggregate coding (Ref. 46)
Rural wages by state × month	Rural wage (₹/day) affordability IV	Labour Bureau WRRI / RBI Handbook wage tables (Refs. 47, 40, safe-scope extension)
HCES MPCE by state × category	Structural category share prior	MoSPI HCES 2022–23 / 2023–24 style (Ref. 48, safe-scope extension)

I.5 Harmonisation steps

1. Retrieve or consult published aggregate releases from the portals in Table I.2 (retrieval date 2026-07-15).
2. Map geography to a canonical state / UT key.
3. Map CPI / public expenditure groups to the six FMCG decision categories (Data chapter).
4. Align annual series (e.g., GSDP) and monthly series (e.g., CPI / TRAI / festivals) onto a common year_month grid.
5. Build feature layers (expected demand, unmet demand, affordability, channel readiness) for OI and the propensity model.

Results and findings use the panel built above. Individual cells need not match one published MoSPI Excel cell.

I.6 Ethics and citation

1. Cite MoSPI, RBI, TRAI, Census, NFHS/UIDAI, india.gov.in, ASCI, Labour Bureau WRRI, and MoSPI HCES via References 37–48.
2. Dataset designation terms are fixed in Table I.1 and the Data chapter. Do not invent alternate labels in results tables.
3. Give the pull date (2026-07-15) when describing data collection / assembly.
4. Claim byte-identical portal dumps only when those files have been newly retrieved and logged here with an updated date.
5. Keep brand customer PII out of every panel.

I.7 CPI lag design

Why cpi_lag1, cpi_lag3, and cpi_yoy? Methodology and Results state the design briefly. Nested holdout tests and diagnostic figures sit here.

I.7.1 Design rationale

Autoregressive CPI terms enter the propensity layer as pre-set classifier features. They match monthly, quarterly, and year-on-year planning horizons.

Table I.5. CPI lag design features — construction and justification.

Feature	Construction	Justification
cpi_lag1	CPI level, (t-1)	Short-run persistence (monthly memory)
cpi_lag3	CPI level, (t-3)	Quarter memory, matches quarterly / pre-festival planning cadence
cpi_yoy	12-month % change	Year-on-year inflation comparison plus seasonal horizon

Prediction target: probability that next-month CPI growth exceeds the national median that month (relative momentum).

I.7.2 Lag-spec comparison

Nested holdout specs (Random Forest reference, same 3-month time holdout, non-CPI covariates fixed) compare richer CPI lag sets step by step:

Table I.6. Nested holdout lag-spec comparison (ROC-AUC and average precision).

Spec	CPI features	Holdout ROC-AUC	Holdout AP
A	none	0.518	0.539
B	lag1	0.516	0.531
C	lag1 + lag3	0.480	0.517
D (baseline)	lag1 + lag3 + yoy	0.736	0.724
E	D + lag6 + lag12	0.737	0.731

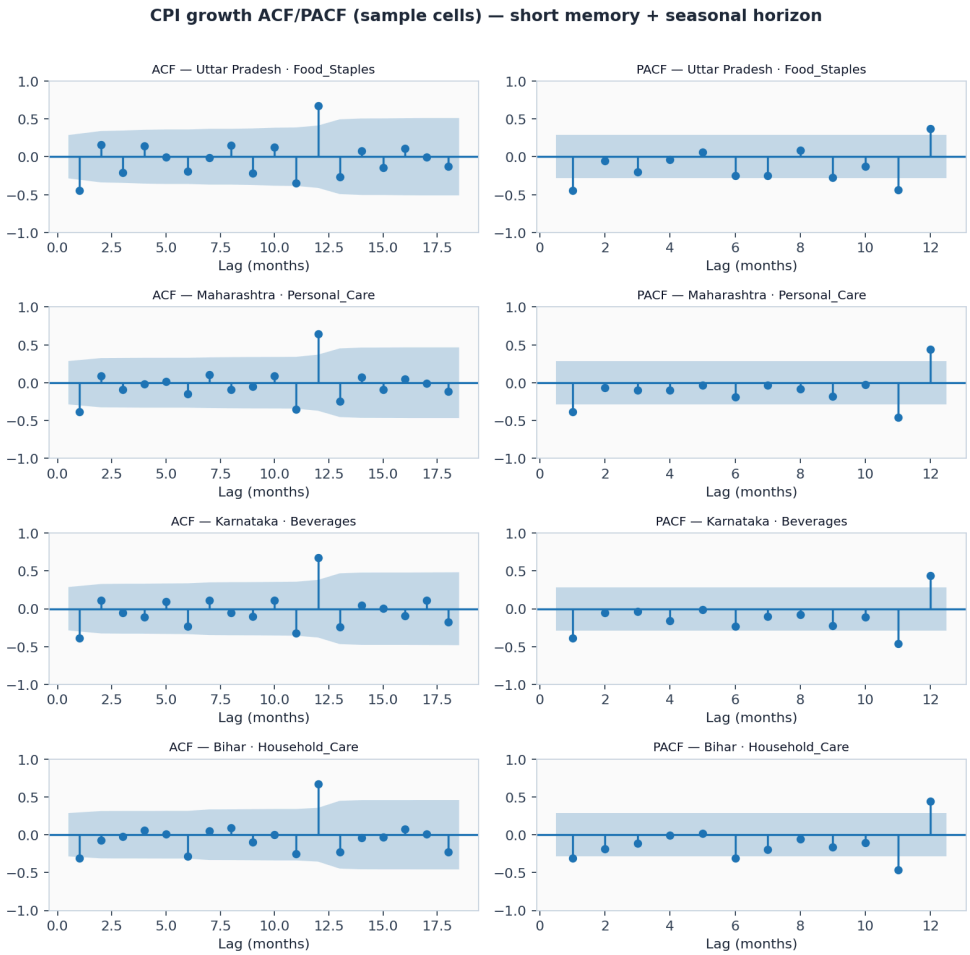
Reading: cpi_yoy is the key CPI feature on this panel. Spec A leads B–C (lags without YoY). Baseline D matches the production feature set and ties richer E. D is kept for simplicity.

I.7.3 Stationarity note

On sample state×category series, CPI levels often fail a unit-root test. Growth series pass. Growth (and the relative-momentum label) is the right object for propensity modelling. Level ARIMA is not the main Decision Stack engine.

I.7.4 Diagnostic figures

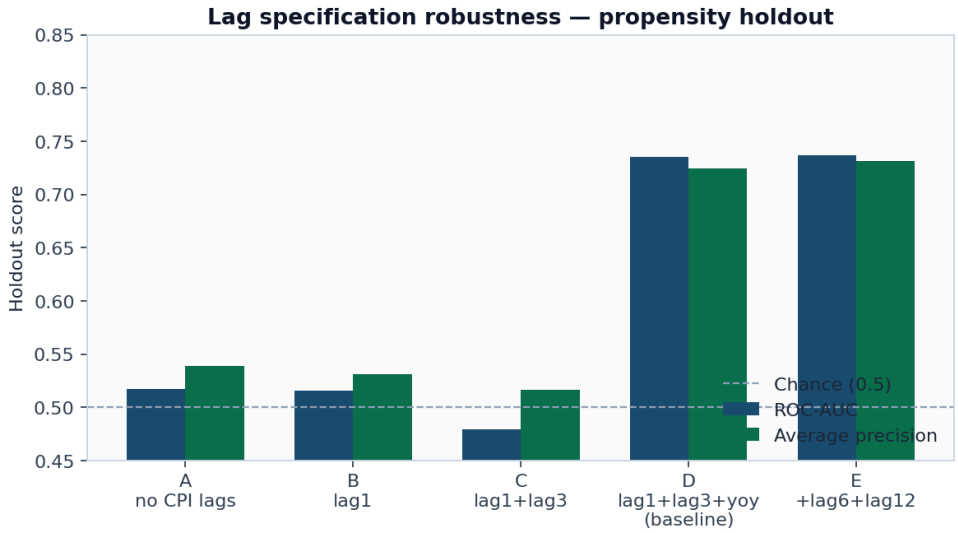
CPI growth ACF and PACF for sample cells



CPI growth ACF and PACF for sample cells

Fig. I.1. Sample-cell ACF/PACF of CPI growth — visual support for short memory and a longer seasonal horizon.

Nested lag-spec holdout ROC-AUC



Nested lag-spec holdout ROC-AUC

Fig. I.2. Nested lag-spec holdout AUC (specs A–E). Baseline D (lag1 + lag3 + yoy) is preferred over A and is close to E.

I.8 Defence diagnostics (axis, weights, stress tests)

Exam checks regenerate with:

```
python scripts/defence_diagnostics.py
```

Files land under data/processed/: oi_propensity_correlation.csv, oi_weight_tilt.csv, oi_component_vif.csv, priority_split_agreement.csv, festival_flag_coverage.csv, sensitivity_discriminative.csv. Headline numbers for the 2025-12 snapshot used in Results:

Table I.7. Defence diagnostics summary (latest-month surface).

Check	Result on reported snapshot
OI ↔ growth propensity (Pearson / Spearman)	0.086 / 0.199
Equal vs affordability-heavy OI (Pearson / Spearman)	0.948 / 0.922
OI-component VIF (max)	4.45 (z_channel_readiness)
Median vs 60th-percentile priority agreement	81.2% (Invest 55 → 37)
Festival flag share on latest month	1.00 (festival-only filter uninformative that month)
Fairness floor pool 20% / 40% / 60% Invest (₹ crore)	3.38 / 3.23 / 3.06

Full scenario rows for varied stress tests match Table 6.10 in Results and findings.

I.8.1 Tier-2 cross-check artefacts

The same script also writes:

1. tier2_continuous_vs_classifier.csv
2. tier2_calibration_reliability.csv
3. tier2_holdout_length_bakeoff.csv
4. tier2_walkforward_winners.csv
5. tier2_calendar_2025_holdout.csv

Table I.8. Tier-2 cross-check summary (matches Tables 6.5–6.6).

Check	Result
Ridge continuous rank-AUC / AP	0.603 / 0.610
Logit classifier AUC / AP (same holdout)	0.759 / 0.748
Spearman(cont, logit), top-quartile overlap	0.415, 48%
Brier (logit holdout probabilities)	0.207
Bake-off winner at 3- / 6-month holdout	logit / logit (0.761 / 0.749)
Walk-forward ends Aug–Nov 2025	logit wins all four windows
Calendar holdout train≤2024-12 / test 2025	AUC 0.734, AP 0.729, n_train 4,464, n_test 2,046

Tier-3 scope note. Industrial brand offtake back-tests and district / city / PIN pilots are not reported. The working panel is state × category × month through 2025-12. No proprietary scanner or district CPI–GSDP–TRAI join sits in data/raw/. Those checks stay future work when data access opens.

DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-42

PLAGIARISM VERIFICATION

Title of the Dissertation: "Designing a Market Decision Stack for Fair Budget Allocation in India FMCG"

Total Pages: 91	Name of the Student: Satyendra
Supervisor: Dr. Deepti Aggrawal	Department: USME

This is to report that the above dissertation was scanned for similarity detection. Process and outcome is given below:

Software used: TURNITIN	Similarity Index: 3%
Word Count: 14244	

Date: 31/07/2026

Candidate's Signature: Satyendra Supervisor(s) Signature: Deepti

Satyendra Satyendra

Project final report

Major Research Project

Document Details

Submission ID
trn:oid:::27535:147023783

Submission Date
Jul 31, 2026, 10:29 PM GMT+5:30

Download Date
Jul 31, 2026, 10:37 PM GMT+5:30

File Name
Project final report .pdf

File Size
3.9 MB

91 Pages

24,458 Words

136,362 Characters

3% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Filtered from the Report

- Bibliography
- Quoted Text
- Small Matches (less than 10 words)

Match Groups

26 Not Cited or Quoted 3%
Matches with neither in-text citation nor quotation marks

0 Missing Quotations 0%
Matches that are still very similar to source material

0 Missing Citation 0%
Matches that have quotation marks, but no in-text citation

0 Cited and Quoted 0%
Matches with in-text citation present, but no quotation marks

Top Sources

2% Internet sources
1% Publications
2% Submitted works (Student Papers)

Match Groups

- **26** Not Cited or Quoted 3%
  Matches with neither in-text citation nor quotation marks
- **0** Missing Quotations 0%
  Matches that are still very similar to source material
- **0** Missing Citation 0%
  Matches that have quotation marks, but no in-text citation
- **0** Cited and Quoted 0%
  Matches with in-text citation present, but no quotation marks

Top Sources

2% Internet sources
1% Publications
2% Submitted works (Student Papers)

Top Sources

The sources with the highest number of matches within the submission. Overlapping sources will not be displayed.

1	Student papers	Delhi Technological University on 2026-06-20 <1%
2	Student papers	Delhi Technological University on 2026-06-12 <1%
3	Internet	docslib.org <1%
4	Student papers	Delhi Technological University on 2026-06-30 <1%
5	Student papers	Delhi Technological University on 2026-06-06 <1%
6	Internet	www.orfonline.org <1%
7	Internet	dtu.ac.in <1%
8	Student papers	Delhi Technological University on 2026-05-31 <1%
9	Internet	www.kse.org.ua <1%
10	Student papers	Delhi Technological University on 2026-05-29 <1%

11	Internet	
www.coursehero.com		<1%
12	Internet	
univendspace.univen.ac.za		<1%
13	Internet	
www.adb.org		<1%
14	Publication	
James Pick, Avijit Sarkar. "33Chapter 2 Conceptual Models of Digital Divides", Wal...		<1%
15	Internet	
assets-eu.researchsquare.com		<1%
16	Internet	
virtusinterpress.org		<1%
17	Internet	
www.ajist.org		<1%
18	Student papers	
Alliance University on 2026-01-30		<1%
19	Student papers	
Delhi Technological University on 2026-05-07		<1%
20	Internet	
open.library.ubc.ca		<1%
21	Internet	
www.ijariit.com		<1%
22	Internet	
www.mdpi.com		<1%