sdmimd
Shri Dharmasthala Manjunatheshwara Institute for Management Development, Mysuru, India
11th International Economics Conference on Economic Growth and Sustainable Development: Innovation, Inclusion and Resilience – November 26-27, 2026

From Price Support to Fiscal Risk: A Machine-Learning Framework for Estimating India's Contingent Liability under the Minimum Support Price System

Sai Prabhakar Balantrapu
Assistant Professor & Doctoral Research Scholar

G. Raghavender Raju
Dean of Humanities & Social Sciences and Language & Literature

Lalith Aditya Kalur
Visiting Faculty
Center of Excellence in Actuarial Data Science
SSSIHL, Bengaluru, Karnataka, India
saiprabhakarbalantrapu@sssihl.edu.in

Abstract

India's Minimum Support Price (MSP) system anchors price support and public procurement for more than twenty crops, yet assessments of its fiscal burden remain almost entirely retrospective, drawn from realised expenditure reported after the fact in budget documents. This paper develops and empirically tests a predictive-analytics framework that forecasts monthly crop prices and market arrivals and converts these forecasts into a forward-looking estimate of the price-support exposure, the contingent fiscal liability, that the MSP creates for the government. Applying the framework to wheat and maize, with eighteen years (April 2008–March 2026) of AGMARKNET mandi-level data aggregated to national monthly series and eight models (ARIMA, SARIMA, Ridge regression, Support Vector Regression, Random Forest, Extra Trees, Gradient Boosting and XGBoost) evaluated one month ahead on a 39-month held-out window (January 2023–March 2026), the study finds that linear time-series models dominate price forecasting: SARIMA is the best model for both crops (MAPE 1.77% for wheat and 2.29% for maize), and no non-linear model is best for any of the twenty-two MSP crops in the wider benchmark. The two crops differ sharply in fiscal profile. Wheat traded below its MSP in 13.7% of months between April 2009 and March 2026 and in a single month of the test window, with a realised liability of ₹174 crore that the estimate flagged correctly. Maize traded below its MSP in 40.7% of months and in 13 of the 39 test months, with a realised liability of ₹5,295 crore, about thirty times wheat's; the estimate captured ₹3,438 crore of it, an under-estimate of 35%, because arrivals forecasts failed to anticipate the size of the most recent harvests. The liability measured is an exposure on the AGMARKNET reporting-market basis, not procurement expenditure. The paper concludes with policy proposals for embedding forecast-based, crop-disaggregated fiscal-risk disclosure in India's MSP budgeting architecture, drawing on international contingent-liability practice.

Keywords: Minimum Support Price; Contingent Fiscal Liability; Machine Learning; Agricultural Price Forecasting; Public Finance; Predictive Analytics; India

Introduction

India's Minimum Support Price (MSP) system, administered through the Commission for Agricultural Costs and Prices (CACP) and operationalised through the Food Corporation of India, state procurement agencies, and, since 2018, the umbrella Pradhan Mantri Annadata Aay Sanrakshan Abhiyan (PM-AASHA), extends price assurance to more than twenty crops spanning cereals, pulses, oilseeds, and commercial crops. The instrument was conceived with a twin objective: to guarantee farmers a floor price that insulates them from the volatility inherent in agricultural markets, and to build the public foodgrain stocks that underwrite India's food-security architecture. Over six decades, MSP has evolved from a narrowly targeted wheat- and rice-procurement mechanism into a wide-ranging price-support regime whose fiscal footprint is substantial and growing: paddy procurement rose from 4,590 lakh tonnes over 2004–05 to 2013–14 to 7,608 lakh tonnes over 2014–15 to 2024–25, and the MSP value paid to farmers for the fourteen kharif crops rose from ₹4.75 lakh crore to ₹16.35 lakh crore across the same two periods (Press Information Bureau [PIB], 2025a); in the 2024–25 rabi marketing season alone, the Food Corporation of India procured 266 lakh tonnes of wheat (PIB, 2025b); and the Union Government's food subsidy is budgeted at ₹2,27,629 crore for 2026–27 (PRS Legislative Research, 2026).

Despite this scale, the fiscal assessment of MSP operations in India remains overwhelmingly retrospective. Procurement expenditure is reported ex post in budget documents, press releases, and parliamentary disclosures, but no systematic, forward-looking, crop-disaggregated method exists for anticipating the government's exposure under evolving price and market-arrival conditions. This is a consequential gap. The international public-finance literature has long argued that price-support and procurement commitments of this kind constitute contingent liabilities — obligations that materialise only if a specified condition occurs — and that treating such liabilities as a residual, ex post accounting item rather than a forecastable exposure leaves fiscal planning structurally reactive (Polackova, 1999; Cebotari, 2008; Demaestri & Moskovits, 2015; International Monetary Fund [IMF], 2016). India's MSP system, precisely because its trigger (market price falling below the notified floor) is itself forecastable from publicly available market data, is an unusually strong candidate for a predictive rather than a purely retrospective approach to fiscal risk management.

The case for a predictive approach is sharpened by a structural asymmetry that recurs across MSP-covered crops but has received little systematic empirical attention: different crops relate to their MSP floor in observably different ways. Wheat, the authors' own multi-crop empirical work finds, traded below its MSP in only 13.7% of months between April 2009 and March 2026, and in a single month of the 39-month test window examined here; procurement at the MSP is a standing operation for wheat, and the market price has stayed above the floor in most months, so the government's price-support exposure is comparatively rare. Maize, by contrast, traded below its MSP in 40.7% of months over the same period and in 13 of the 39 test months. Procurement of coarse cereals is conducted on a limited scale, so that a notified floor exists on paper while the gap between the market price and the floor is borne largely by farmers, and becomes a claim on the public purse only if the guarantee is made effective. A crop-agnostic, flat approach to MSP fiscal budgeting cannot capture this heterogeneity, and a national early-warning system built on such an approach risks materially mis-provisioning across the crop portfolio.

This paper develops and empirically demonstrates a predictive-analytics framework that (i) forecasts crop prices and market arrivals one month ahead using a suite of machine-learning and classical time-series models, (ii) combines these forecasts into an estimate of the government's MSP-linked price-

ISBN code 978-81-999803-3-4

support exposure and evaluates it against realised liability through a four-way classification (sensitivity, precision, accuracy) and a rupee-value comparison, and (iii) states precisely what that exposure measure is, and is not, relative to the government's procurement expenditure. The framework is applied to wheat and maize as paired, structurally contrasting case studies, with portfolio-wide results for twenty-two crops used for context. The remainder of the paper proceeds as follows. Section 2 situates the paper within the agricultural price-risk, machine-learning forecasting, and contingent-liability literatures. Section 3 sets out the data, forecasting methodology, and liability framework. Section 4 presents the empirical findings for wheat and maize, places them in the context of the wider benchmark and draws policy implications. Section 5 concludes with limitations and directions for further research.

Review of Literature

The theoretical starting point for this paper is the classical distinction between risk and uncertainty articulated by Knight (1921), in which risk denotes outcomes whose probability distribution is knowable, and uncertainty denotes outcomes that are not. Agricultural output prices sit closer to the risk end of this spectrum than is often assumed: while individual shocks (weather, pest incidence, global price movements) are themselves uncertain, the resulting price series exhibit statistical regularities — seasonality, mean reversion around policy-anchored floors, and serial correlation — that are amenable to formal estimation. Just and Pope (1978, 1979) formalised this insight for agricultural production functions by proposing stochastic specifications that separate the mean and variance effects of input choices, a foundation on which subsequent empirical work on production and price risk has built.

The forecasting methodology employed in this paper draws on two distinct traditions. The classical time-series tradition, exemplified by the Box and Jenkins (1976) ARIMA framework and its seasonal SARIMA extension, models a price series as a function of its own past values and remains a strong benchmark for agricultural commodities whose seasonality is stable and whose autocorrelation structure is close to linear. The statistical-learning tradition offers a complementary set of tools better suited to non-linear and high-dimensional relationships: ridge regression (Hoerl & Kennard, 1970) regularises linear estimation to control variance under multicollinearity, an especially relevant property for lag-feature-heavy agricultural price models; random forests (Breiman, 2001) and gradient boosting (Friedman, 2001), including the widely adopted XGBoost implementation (Chen & Guestrin, 2016), extend prediction to non-linear, interaction-rich relationships; and recurrent architectures such as the long short-term memory network (Hochreiter & Schmidhuber, 1997) extend the toolkit further to sequence-dependent structure, though at a data-volume cost that is often prohibitive for a single-crop, monthly national series of the length available here. A recent systematic review of machine-learning applications to staple-grain price forecasting (Theofilou et al., 2025) reports R^2 values of about 0.86 to 0.96 and a MAPE of about 4% (4.09%) for a Thai rice benchmark, figures against which this paper's own wheat and maize results are benchmarked in Section 4.

A second relevant literature concerns agricultural insurance and revenue protection. Miranda and Glauber (1997) demonstrate that crop insurance markets are exposed to systemic, spatially correlated risk that conventional actuarial pooling cannot diversify away — a finding with direct relevance to India's own Pradhan Mantri Fasal Bima Yojana (PMFBY) and to the broader revenue-protection agenda of which the present paper's fiscal-liability framework forms one component. While revenue-protection design (Value-at-Risk and Conditional Value-at-Risk-based premium pricing, Gaussian copula-based price-volume dependence modelling) lies outside this paper's immediate scope, the

systemic-risk finding underlines why a government-borne contingent liability, rather than a purely market-based insurance instrument, remains a first-order consideration for staple crops whose price risk is nationally correlated.

The third and most directly relevant literature is that on contingent liabilities and fiscal risk management. Polackova (1999) introduced the now-standard typology distinguishing explicit from implicit, and direct from contingent, government liabilities, arguing that implicit contingent liabilities — obligations governments are expected, though not legally bound, to honour — are systematically under-disclosed in conventional budgeting. Cebotari (2008) extends this typology into an operational IMF framework for identifying, quantifying, and disclosing contingent liabilities across a range of country contexts, while Demaestri and Moskovits (2015) and the IMF's (2016) own guidance on fiscal risk statements set out best-practice approaches to monitoring such exposures over time rather than merely recording them after realisation. The World Bank (2005) and OECD (2009) frameworks for agricultural risk management situate price-support commitments explicitly within this contingent-liability universe, a positioning this paper adopts directly: India's MSP system is best understood not as a fixed budget line but as a state-contingent claim whose magnitude depends on the joint realisation of price and market arrivals — precisely the two quantities a predictive-analytics framework is well suited to forecast.

Taken together, these literatures identify a clear gap that this paper addresses. Agricultural forecasting research has developed increasingly sophisticated price-prediction tools, and public-finance research has developed increasingly sophisticated contingent-liability frameworks, but the two have rarely been integrated for India's MSP system specifically. No study known to the authors combines crop-specific machine-learning price and arrivals forecasting with a formal, classification-based contingent-liability estimation exercise evaluated against realised liability. This paper offers such an integration, using wheat and maize as a paired illustration of how structurally different crops generate structurally different fiscal-risk profiles.

Study Objectives and Methodology

Objectives

The paper pursues two linked objectives. The first is to forecast monthly prices and market arrivals for wheat and maize one month ahead using a systematically benchmarked suite of machine-learning and time-series models, and to identify the model family best suited to MSP-anchored agricultural series. The second is to convert these forecasts into a forward-looking estimate of the government's MSP-linked price-support exposure, the contingent liability that arises when the market price falls below the notified MSP, to evaluate its rupee-value error and its performance as an early-warning signal against realised outcomes, and to state clearly how this exposure measure relates to procurement expenditure.

Data

The empirical analysis draws on eighteen years of mandi-level daily price and arrival records for wheat and maize, sourced from AGMARKNET (April 2008–March 2026). The modal price is used as the measure of market price. For each month, the national price is the average of the modal prices of all

valid daily records and national arrivals are the sum of recorded arrivals, after cleaning for duplicates, unit inconsistencies and implausible values. The monthly series are aligned with the MSP of the marketing season then in force, using CACP and Department of Agriculture and Farmers Welfare notifications. The resulting series span 216 months for each crop, split chronologically into a 177-month training window (April 2008–December 2022) and a 39-month held-out test window (January 2023–March 2026), with the split fixed in advance to avoid look-ahead bias. A validation window (January 2021–December 2022), for which the models are re-estimated on data to December 2020, is used for model selection (Section 3.3). Because AGMARKNET covers regulated wholesale markets, the series measure the marketed surplus passing through those markets, which is the quantity relevant to procurement, rather than production, and the national average can diverge from individual markets and states.

Price and Arrivals Forecasting

Eight candidate specifications were benchmarked for both price and arrivals forecasting: ARIMA and SARIMA (with orders selected by AIC on the training data), Ridge regression on lagged values, Support Vector Regression, and four tree ensembles (Random Forest, Extra Trees, Gradient Boosting and XGBoost). Arrivals are modelled as the logarithm of one plus arrivals and transformed back before evaluation. Price forecasts are evaluated on Mean Absolute Error, Root Mean Squared Error, Mean Absolute Percentage Error (MAPE) and the coefficient of determination (R^2), and arrivals forecasts on RMSE expressed as a percentage of mean arrivals and R^2; all are computed one month ahead on the chronologically held-out test window and compared with a no-change forecast. Two rules select a model for each crop. Rule A chooses the model with the lowest test-window error and is a retrospective upper bound on attainable accuracy; rule B chooses the model on the validation window, before the test period is observed, and is the principal basis of evaluation because it reproduces the position of a forecaster who must choose before outcomes are known. For wheat and maize the two rules select the same models. The price and arrivals benchmarks for wheat and maize build on the authors' earlier published studies of these two crops (Balantrapu et al., 2026, in press); the present paper converts the forecasts into a liability estimate and evaluates it against realised liability.

Across the wider 22-crop portfolio of which wheat and maize form part, linear models win the price benchmark for every crop: SARIMA is best for twelve crops (wheat and maize included), ARIMA for five and Ridge regression for five, while neither Support Vector Regression nor any tree ensemble is best for any crop. Three explanations are advanced. The first is the persistence of storable-commodity prices implied by the theory of storage, which favours models built on first differences. The second is the steadily rising MSP, which imparts a persistent upward drift to prices. The third is structural: because Random Forest, Extra Trees, XGBoost and Gradient Boosting predict by averaging training-set leaf values, they cannot extrapolate beyond the range of prices observed in training, a material handicap on a trending test series whose prices frequently exceed anything seen in the training period. The result is a finding about the eight models as specified here (univariate lag features, fixed settings, national monthly series, a one-month horizon and the 2023–26 window), not a claim that linear methods are generally superior to non-linear ones.

Contingent-Liability Classification Framework

Let P(c,t) denote the realised national market price of crop c in month t, MSP(c,t) the notified floor in force, and Q(c,t) market arrivals in the reporting markets. Realised liability is L(c,t) = max(0, MSP(c,t) − P(c,t)) × Q(c,t), expressed in rupees crore, and a month is classified as a liability month when L(c,t) is positive, that is, when the price is below the MSP by any amount, however small. The estimated

liability $\hat{L}(c,t) = \max(0, MSP(c,t) - \hat{P}(c,t)) \times \hat{Q}(c,t)$ is computed identically from the one-month-ahead forecasts of price and arrivals. Comparing the months signalled by the estimate with the months in which liability arose over the 39-month test window yields, for each crop, a four-way classification (true positives, false negatives or missed liability, false alarms and true negatives) from which sensitivity, precision and accuracy are computed in the conventional way; the rupee-value estimation error is the difference between estimated and realised liability as a percentage of the realised value. Both are compared with a no-change benchmark in which each month's liability is computed from the previous month's price and arrivals, so that the marginal contribution of forecasting can be separated from the persistence of below-MSP conditions.

Scope of the Liability Measure

An important clarification is required before the results in Section 4 are interpreted. The liability $L(c,t)$ is an exposure measure. It is the amount the government would have to pay to make good the full shortfall of the market price below the MSP on all produce arriving in the AGMARKNET reporting markets, whether by procurement at the MSP or by a price-deficiency payment, and it is therefore an upper bound on the price-support element of the MSP commitment, not a record of what was spent. It is zero whenever the market price is at or above the MSP, however much the government may procure in that month, and positive whenever the price falls below the MSP, whether the government procures anything. It is not procurement expenditure, a subsidy payment or the realised loss of farmers. The distinction matters most for wheat and rice, for which procurement at the MSP is a standing operation that supplies the Public Distribution System and the buffer stock; the costs of acquiring, storing and distributing the grain are recorded in the food subsidy, depend on the quantities procured and on handling costs, and are not modelled here. The estimates are therefore evaluated against the liability realised in the market, in total, month by month and as an early warning, and not against budgetary outlays. A further qualification concerns aggregation: because the liability responds only to prices below the MSP, applying it to a national average price yields in general a smaller figure than applying it market by market and adding up, so the measure is best read as a conservative estimate of exposure. The size of this understatement is not quantified here.

Analysis and Discussion

Price and Arrivals Forecasting Performance

Table 1 summarises the best-performing model and its test-window accuracy for both forecasting tasks and both crops. SARIMA is the best model for both crops on both tasks, and the validation-selected (rule B) and test-selected (rule A) choices coincide. For wheat the price MAPE is 1.77% (MAE ₹44.4 per quintal, RMSE ₹64.2, R^2 0.848) and for maize 2.29% (MAE ₹47.6, RMSE ₹59.8, R^2 0.892); Ridge regression, the best of the non-time-series models, follows at 2.61% and 2.81% respectively. Both MAPEs lie below the about 4% reported by Theofilou et al. (2025) for a Thai rice benchmark; maize's R^2 lies inside the roughly 0.86–0.96 range of that review, whereas wheat's falls slightly below it. The forecasts also improve on a no-change forecast (MAPE 2.63% for wheat and 2.93% for maize), and the improvement is statistically significant for both crops, although a no-change forecast already achieves a MAPE below 3% for these crops, so that part of the accuracy of any model reflects the persistence of prices. The best model predicts the direction of the next month's price change correctly in 71% of

ISBN code 978-81-999803-3-4

months for wheat and 76% for maize. Arrivals are forecast far less accurately than prices. Wheat arrivals are forecast with an RMSE of 21.2% of mean arrivals (R^2 0.974), among the most accurately forecast of the twenty-two crops; maize arrivals are forecast with an RMSE of 39.6% (R^2 0.679). Both are more accurate than a no-change forecast, which places the harvest peak one month late, but maize's arrivals forecast is the weaker link in its liability estimate, as Section 4.2 shows.

Task	Crop	Best model	RMSE	Error (%)	R^2	No-change error (%)
Price	Wheat	SARIMA	64.2	1.77	0.848	2.63
Price	Maize	SARIMA	59.8	2.29	0.892	2.93
Volume	Wheat	SARIMA	4.94	21.2	0.974	155.3
Volume	Maize	SARIMA	3.15	39.6	0.679	58.0

Table 1. Best-Performing Forecasting Model, Wheat and Maize (39-Month Test Window, One-Month-Ahead). RMSE: ₹ per quintal for price, million quintals for volume. Error (%): MAPE for price, RMSE as % of mean arrivals for volume.

The Wheat–Maize Fiscal Asymmetry

The paper's central empirical finding emerges from Table 2, which juxtaposes the two crops' liability results. Wheat's market price stayed above its MSP in 86.3% of months between April 2009 and March 2026, and in 38 of the 39 test months. The single liability month of the test window produced a realised liability of ₹174 crore, which the estimate flagged correctly (₹128 crore estimated, 26.4% below the realised value); wheat's classification scores are therefore perfect, but they rest on one observation and say little about the framework's ability to anticipate a rare event. Maize presents a different profile. Its price was below the MSP in 40.7% of months over the same period and in 13 of the 39 test months (33.3%), and its realised test-window liability was ₹5,295 crore, about thirty times wheat's. The estimate signalled twelve of the thirteen liability months, with two false alarms (sensitivity 92.3%, precision 85.7%, accuracy 92.3%), but totalled ₹3,438 crore, an under-estimate of 35.1%. The under-estimate is a volume error rather than a price error: the arrivals forecasts failed to anticipate the size of the 2024–25 and 2025–26 harvests.

Metric	Wheat	Maize
Liability months in test window	1 of 39 (2.6%)	13 of 39 (33.3%)
Months below MSP, April 2009–March 2026	13.7%	40.7%
True positives / false negatives / false alarms / true negatives	1 / 0 / 0 / 38	12 / 1 / 2 / 24
Sensitivity	100.0%	92.3%
Precision	100.0%	85.7%
Accuracy	100.0%	92.3%
Realised liability (₹ crore, 39-month window)	174	5,295

ISBN code 978-81-999803-3-4

Metric	Wheat	Maize
Estimated liability (₹ crore)	128	3,438
Estimation error	−26.4%	−35.1%
Cumulative realised liability, 2009–10 to 2025–26 (₹ crore)	3,455	10,519
Share of 22-crop cumulative liability	3.8%	11.7%

Table 2. Liability Classification and Fiscal Impact, Wheat and Maize (39-Month Test Window; Rule A, Which Gives the Same Estimates as Rule B for Both Crops)

The fiscal implication of this contrast is the paper's sharpest finding. Over the seventeen financial years 2009–10 to 2025–26, maize accounts for a cumulative realised liability of ₹10,519 crore (11.7% of the ₹90,100 crore cumulative total for twenty-two crops) against ₹3,455 crore (3.8%) for wheat, and in the test window maize contributes 15% of the portfolio's realised liability against 0.5% for wheat, although wheat, with paddy, accounts for about half of the market value of arrivals of the twenty-two crops. A budgeting rule that allocated contingent provision in proportion to crop economic size would therefore over-provide for wheat and under-provide for maize, whose price is more often below the floor and whose procurement is limited. The finding also carries a caution about estimation: the crop with the larger exposure is the crop whose estimate is least accurate, because its liability is driven by harvest-time arrivals, the hardest component to forecast, whereas wheat's exposure depends on small movements of the price around the MSP. High classification scores should be read with care, because below-MSP conditions are persistent: across all 858 crop-months of the twenty-two crops, a no-change rule also scores highly (about 92%), and the contribution of forecasting lies in the turning points, where the estimate correctly classifies 41% of the months in which a crop moves into or out of liability against 9% for the no-change rule.

Portfolio Context and the Distinction between Exposure and Expenditure

Table 3 places the two crops in the context of the full twenty-two-crop benchmark. Over the test window the realised liability of the twenty-two crops totalled ₹34,697 crore, arising in 462 of 858 crop-months. With models chosen in advance (rule B), the estimate totals ₹32,269 crore, an under-estimate of 7.0%; with the test-selected models (rule A) the total is within 0.1% of the realised value, but this closer figure reflects hindsight in model choice. The totals conceal offsetting errors: summed crop by crop and month by month, the absolute errors amount to about 40% of the realised liability, and a no-change benchmark comes closer to the cumulative total (+1.2%) because its errors of timing cancel when months are added. The value of forecasting lies in timing: the mean absolute error of the monthly aggregate estimate is about 20% of the realised monthly liability against 41% for the no-change rule, and its correlation with the realised monthly series is 0.96 against 0.81.

Measure	Rule B (principal)	Rule A (upper bound)
Estimated total vs realised ₹34,697 crore	₹32,269 crore (−7.0%)	₹34,648 crore (−0.1%)

ISBN code 978-81-999803-3-4

Measure	Rule B (principal)	Rule A (upper bound)
Sensitivity	93.3%	94.4%
Precision	92.7%	92.8%
Accuracy	92.4%	93.0%

Table 3. Early-Warning Performance and Total Estimate, Twenty-Two MSP Crops, 858 Crop-Months (39-Month Test Window)

The liability measured here should not be confused with the cost of the MSP system. The food subsidy was budgeted at ₹2,27,629 crore for 2026–27 (PRS Legislative Research, 2026) and is dominated by the normal-course procurement of paddy and wheat at and above the MSP, operations that are the policy itself rather than a liability in the sense defined here, and by the costs of storing and distributing grain. The estimated exposure isolates a different quantity: the shortfall of market prices below the support price on produce reaching the reporting markets, which was generally below ₹2,000 crore a year until 2016–17, settled at ₹8,000–11,000 crore a year from 2017–18 and reached a record ₹20,800 crore in 2025–26, to which maize contributed ₹5,100 crore. The two are not alternative measures of the same outlay, and it is precisely because this exposure is neither budgeted nor disclosed in advance that its estimation is of interest. Combining the exposure with the cost of procurement would require crop-level procurement quantities and expenditure, which this paper does not model.

Policy Implications and International Comparators

Four policy directions follow from these findings; they are proposals for evaluation and piloting rather than validated designs. First, the monthly crop-wise liability estimate could be published alongside the Union Budget as a statement of estimated contingent exposure, with a margin for forecast error derived from back-testing (about 20% for the monthly total and 7% over the test window), in keeping with the fiscal-risk-disclosure practice recommended by Polackova (1999), Cebotari (2008) and the IMF (2016). Second, instruments could be differentiated by crop: for wheat-type crops, where procurement is a standing operation, monitoring should focus on the proximity of the market price to the floor, since small price errors are magnified there, while for maize-type crops, where procurement is limited, a forecast-triggered price-deficiency payment at harvest merits a pilot; the cost of any such scheme depends on its reference price, threshold and coverage and cannot be read directly from the liability estimate. Third, because 59% of the test-window liability arose in the four months from October to January, provision could be made seasonally and backed by a counter-cyclical stabilisation reserve built up in years of low liability, along the lines of the existing Price Stabilisation Fund. Fourth, the exposure could be estimated at the level of markets and states and combined with crop-level procurement quantities and costs, to show how much of it is met through procurement and how much falls on farmers. These proposals find international analogues: the United States' Agriculture Risk Coverage programme pays when county-level revenue falls below a benchmark, and the European Union's Income Stabilisation Tool compensates farmers for severe falls in income and is triggered only by large losses. Both offer India-relevant precedent for building forecastability into a price-support instrument's fiscal architecture rather than treating fiscal risk as an unavoidable by-product of the policy.

ISBN code 978-81-999803-3-4

Conclusion and Scope for Further Study

This paper has developed and empirically tested a predictive-analytics framework for estimating India's MSP-linked contingent fiscal exposure, applying it to wheat and maize as paired, structurally contrasting case studies. Three findings stand out. First, linear time-series models dominate price forecasting: SARIMA is the best model for both crops (MAPE 1.77% for wheat and 2.29% for maize), and across the twenty-two-crop benchmark no non-linear model is best for any crop, a result grounded in the persistence of prices, the upward drift imparted by the rising MSP and tree ensembles' inability to extrapolate beyond their training range. Second, wheat and maize present sharply different fiscal profiles: maize's realised test-window liability of ₹5,295 crore is about thirty times wheat's ₹174 crore, and maize traded below its MSP in 40.7% of months against 13.7% for wheat, so that provision proportional to crop economic size would under-provide for maize; yet maize is also the crop whose exposure is estimated least accurately (an under-estimate of 35.1%), because its arrivals are hard to forecast. Third, across all twenty-two crops the exposure can be estimated one month ahead from public data to within 7% of the realised total when models are chosen in advance, and the estimate signals 93% of liability months, but the contribution of forecasting lies in the timing of liability and not in the cumulative total, where a no-change rule comes closer.

The paper's principal limitations define its clearest directions for further work. First, the estimate looks one month ahead over a single 39-month test window that contains distinctive episodes, including the commodity-price shock that followed the conflict in Ukraine and the fall in maize prices in 2025–26; the accuracy reported is that of a monthly monitoring tool, not of an annual budget forecast, and longer horizons would be less accurate. Second, the liability is formed from point forecasts of price and arrivals, without allowance for the correlation of their errors and without a prediction interval; wheat's rare, narrowly triggered liability in particular would be better treated with a Monte Carlo or interval-based approach around the MSP threshold. Third, the measure is an exposure computed on a national average price and the reporting-market basis of AGMARKNET; it understates what a market- or state-level calculation would show and is not procurement expenditure. A natural extension is to combine this price-support exposure with a procurement-cost component for the procured share of arrivals, unified through an empirically estimated procurement-coverage ratio that prevents double-counting between procured and non-procured shares; estimating that ratio would require crop-level procurement-quantity data, which this paper does not use. The wheat–maize comparison also forms part of a wider doctoral research programme covering all twenty-two MSP crops, whose natural extension is a national early-warning dashboard aggregating crop-specific exposure forecasts into a single policy-actionable instrument.

More broadly, this paper argues that India's MSP system is best understood, and best budgeted for, as a state-contingent fiscal claim rather than a fixed line item — a reframing consistent with established international contingent-liability practice and directly relevant to the broader project of building resilient public finances in an era of climate variability, volatile global commodity markets, and shifting trade relationships. Embedding predictive analytics into India's fiscal-risk architecture for agricultural price support offers a concrete, empirically grounded step toward that resilience, one this paper has sought to demonstrate is both methodologically tractable and, as the wheat–maize asymmetry shows, policy-relevant in ways a purely retrospective accounting approach would not have revealed.

References

Balantrapu, S. P., Raghavender Raju, G., Mudigonda, S. S., & Kalur, L. A. (2026). Forecasting Minimum Support Price-anchored mandi prices for wheat and maize in India: A comparative machine-learning

and time-series framework. Agricultural Science Digest, 46, 129–135. https://doi.org/10.18805/ag.D-6567

Balantrapu, S. P., Raghavender Raju, G., Mudigonda, S. S., & Kalur, L. A. (in press). Comparative analysis of machine learning and time series models for crop volume prediction: Evidence from AGMARKNET data for wheat and maize in India. The Indian Journal of Agricultural Sciences.

Box, G. E. P., & Jenkins, G. M. (1976). Time series analysis: Forecasting and control (Rev. ed.). Holden-Day.

Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32.

Cebotari, A. (2008). Contingent liabilities: Issues and practice (IMF Working Paper WP/08/245). International Monetary Fund.

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785–794). Association for Computing Machinery.

Demaestri, E., & Moskovits, C. (2015). Contingent liabilities: A threat to fiscal sustainability. Inter-American Development Bank.

Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. Annals of Statistics, 29(5), 1189–1232.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.

Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. Technometrics, 12(1), 55–67.

International Monetary Fund. (2016). Analyzing and managing fiscal risks: Best practices (IMF Policy Paper). International Monetary Fund.

Just, R. E., & Pope, R. D. (1978). Stochastic specification of production functions and economic implications. Journal of Econometrics, 7(1), 67–86.

Just, R. E., & Pope, R. D. (1979). Production function estimation and related risk considerations. American Journal of Agricultural Economics, 61(2), 276–284.

Knight, F. H. (1921). Risk, uncertainty and profit. Houghton Mifflin.

Miranda, M. J., & Glauber, J. W. (1997). Systemic risk in agricultural crop insurance markets. American Journal of Agricultural Economics, 79(1), 206–215.

Organisation for Economic Co-operation and Development. (2009). Managing risk in agriculture: A holistic approach. OECD Publishing.

Polackova, H. (1999). Contingent government liabilities: A hidden risk for fiscal stability (World Bank Policy Research Working Paper No. 1989). World Bank.

Press Information Bureau. (2025a, May 28). Cabinet approves minimum support prices (MSP) for kharif crops for marketing season 2025-26 [Press release]. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2131983

ISBN code 978-81-999803-3-4

Press Information Bureau. (2025b, October 10). Minimum support prices: From safety net to self-sufficiency [Backgrounder]. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2177219

PRS Legislative Research. (2026, February 1). Union budget 2026-27 analysis. https://prsindia.org/files/budget/budget_parliament/2026/Union_Budget_Analysis-2026-27.pdf

Theofilou, A., Nastis, S. A., Michailidis, A., Bournaris, T., & Mattas, K. (2025). Predicting prices of staple crops using machine learning: A systematic review of studies on wheat, corn, and rice. Sustainability, 17(12), Article 5456. https://doi.org/10.3390/su17125456

World Bank. (2005). Managing agricultural production risk: Innovations in developing countries. World Bank.

ISBN code 978-81-999803-3-4