

Emergency Cash Buffers and Investment Outcomes during Income Shocks in India

Harini R¹, Srihari V², Gopavarapu Sreekar³

^{1,2,3}Department of Finance, Universal AI University, Karjat

Abstract

This study examines how emergency-reserve size and withdrawal order affect equity liquidation, reserve recovery and investment-goal affordability for synthetic asset-owning Indian households. A deterministic historical simulation combines monthly Nifty 50 Total Return Index and adjusted HDFC Liquid Fund Direct Growth returns for January 2014 to December 2025. Forty-five scenario definitions, seven reserve sizes and 85 overlapping five-year windows produce 26,775 conditional paths. Household extensions use published urban spending and earnings benchmarks while explicitly varying saving capacity. For the original household facing full income loss during months 12 to 17, a six-month reserve reduces mean spending-funded equity sales from Rs250,000 to Rs5,777, but lowers mean terminal wealth from Rs2,394,214 to Rs2,278,077. Seven months is the smallest tested reserve that eliminates equity sales in every window for this scenario. The six-month policy increases the mean of the lowest 10% of wealth outcomes by Rs35,445 relative to zero reserve. Yet its goal-success share falls from 12.9% to 0% when a common terminal reserve must also be retained. No path with a 5% or 10% initial saving share rebuilds the six-month reserve within the horizon after the full interruption. These findings distinguish immediate spending protection from long-term affordability and recovery capacity. They support conditional policy comparisons rather than a universal reserve recommendation or estimated future success probability.

Keywords: Emergency Savings; Household Finance; Income Shocks; Liquidity Reserves; Goal Affordability; India; Historical Simulation

1. Introduction

Households that own financial investments face an allocation decision between readily available reserves and assets intended to finance longer-term goals. An income interruption can convert that allocation decision into a withdrawal decision: planned spending must be funded from surviving income, liquid assets or investment sales. A larger reserve may preserve equity holdings through the interruption, but the initial allocation and later replenishment also affect the wealth available for a financial goal. Evaluating an emergency fund therefore requires attention to both the spending it supports and the investment opportunities it displaces.

Precautionary saving and the interaction between labour-income risk and portfolio choice have established theoretical foundations [1, 2]. Practical simulation research has also examined the cost of maintaining separate emergency reserves [3]. This paper takes those mechanisms as its starting point. Its purpose is to make a bounded, reproducible comparison using Indian market histories and transparent household cash

flows, rather than derive a utility-maximising allocation. The distinction matters because avoiding an equity sale, retaining sufficient assets for a goal and recovering a depleted reserve are different objectives. Three measurement problems motivate the analysis. First, a household can finance its current spending by selling equity and still finish with more wealth than a household holding a larger reserve. A zero-sale policy is therefore not automatically the policy with the highest terminal wealth. Second, using total terminal assets to assess a financial goal implicitly allows emergency money to be spent on that goal. A household that must retain a reserve faces a larger terminal obligation. Third, a reserve that covers one interruption may remain depleted when another interruption arrives. The capacity to replenish it depends on the post-shock budget as well as its initial size.

The research question is: for synthetic urban Indian households that already own marketable financial assets, how do reserve size and deficit-withdrawal order affect spending coverage, five-year goal affordability and reserve recovery under observed Indian market-return sequences and specified income interruptions? The experiment compares seven affordable initial reserve sizes within matched market windows. It retains an original higher-saving household and adds profiles with lower saving margins, surviving earnings and an ongoing budget deficit. The added profiles use published consumption and earnings statistics to make their scale explicit, without treating them as sampled households.

The contribution is an integrated set of conditional outcomes. Spending-funded equity sales and unmet planned spending are reported alongside mean and downside terminal wealth. Goal affordability is evaluated with a common reserve retained across all strategies. Recovery is reported separately, including failure to restore the reserve before the horizon ends. Identical-start withdrawal comparisons distinguish the effect of using cash first from the effect of initially allocating more money to cash. Together, these measures expose trade-offs that a single goal-success or zero-sale statistic would conceal.

The main result is that liquidity protection does not establish investment-goal sufficiency. In the original six-month interruption, a six-month reserve substantially reduces equity sales and improves the lower tail of ending wealth, while reducing mean ending wealth. A common retained-reserve requirement makes the financial goal more demanding. Lower-saving profiles can meet planned spending during the interruption yet fail to restore their reserves within five years. The interpretation is conditional on the available return history, assumed initial assets and specified cash-flow rules.

2. Literature and research contribution

2.1 Precautionary saving and household portfolio choice

Carroll [1] links buffer-stock saving to income uncertainty, providing a theoretical reason for households to hold wealth against adverse cash flows. Cocco, Gomes and Maenhout [2] study consumption and portfolio choice with labour income and borrowing constraints. These contributions establish that saving, income risk and asset allocation are related decisions. The present experiment does not solve the corresponding optimisation problem: consumption paths, income interruptions and reserve rules are imposed, and neither risk preferences nor household welfare are estimated.

Scott and colleagues [3] provide the closest methodological comparator. Their US simulation considers emergency adequacy and retirement accumulation. Its traditional reserve strategy uses cash before investment assets and replenishes cash before restarting investment contributions. These rules therefore precede the present study. The extension here uses Indian monthly return sequences, a five-year horizon and jointly reported liquidation, recovery and goal outcomes. Differences in horizon, return construction and household assumptions prevent an exact replication claim.

2.2 Consumption smoothing and recovery after disruption

Indian evidence already documents asset-based consumption smoothing. Imai and Malaeb [4], reviewed through their revised author working paper, examine rural households adjusting currency, crop inventories and capital assets. Malaeb, Golan and Imai [5] further examine how coping mechanisms vary with asset holdings. These studies make geography alone an insufficient novelty claim. The households considered here instead begin with a specified stock of marketable financial assets and allocate it between an equity benchmark and a liquid-fund proxy. Their outcomes cannot be generalised to rural households relying on productive or informal assets.

Farrell, Greig and Yu [6] analyse cash adequacy against adverse household cash flows using US account data. Ganong and Noel [7], reviewed through the working-paper version associated with their journal article, examine consumption during unemployment and the role of benefit exhaustion. Both motivate attention to the timing of resources and spending. Their empirical estimates are not imported as Indian job-loss probabilities. Gupta, Malani and Woda [9] document heterogeneous income and consumption changes during the COVID shock in India, supporting heterogeneous scenarios while not calibrating the particular interruptions used here.

Reserve recovery is also an established concern. Martin and colleagues [8] distinguish the depletion of savings during a crisis from a subsequent period of rebuilding. The current study applies that distinction to an investment setting, asking whether reserve restoration occurs alongside goal accumulation and whether repeated interruptions arrive before restoration. Recovery duration is measured against a moving spending-based target, and incomplete recovery is reported as censoring rather than a zero-duration observation.

2.3 Portfolio benchmarks and the bounded research gap

Estrada [10] evaluates retirement cash-bucket strategies against rebalanced alternatives, illustrating the importance of the comparison policy. A result against an all-equity initial allocation does not establish superiority over every portfolio strategy. Accordingly, this paper includes proportional deficit withdrawals from identical starting assets as a separate benchmark. It does not test periodic constant-mix rebalancing or an optimal dynamic allocation.

The gap addressed is an applied measurement and comparison problem: how do emergency-reserve policies perform when spending coverage, investment liquidation, terminal goal obligations and reserve rebuilding are assessed together under matched Indian return histories? The common terminal reserve makes goal obligations comparable, while saving-margin and repeated-shock scenarios distinguish initial coverage from recovery capacity. These choices extend the reviewed work within a specific design. They are not evidence that each mechanism, or their general combination, has never been studied. The literature review is targeted rather than systematic; theoretical background claims are limited to the verified abstracts and empirical claims to the reviewed primary articles, reports or identified working-paper versions.

3. Data and household construction

3.1 Market series and cleaning

The equity proxy is the Nifty 50 Total Return Index from Nifty Indices [13]. The liquidity proxy is HDFC Liquid Fund Direct Plan Growth, AMFI scheme code 119091, obtained from official historical NAV records [14]. The sample contains 145 matched monthly levels from December 2013 through December 2025, giving 144 return pairs from January 2014 through December 2025. The source archive was

collected on 29 September 2026 and deliberately ends in December 2025; no later market observations are used.

Each monthly observation uses the last date in that calendar month for which both series are available. The liquid-fund NAV requires a common-unit adjustment for the face-value change effective 30 August 2015, documented in the HDFC addendum [15]. Earlier NAV observations are multiplied by 100 and later observations by 1. These multipliers are adjustment factors, not returns. Monthly simple returns equal the current adjusted level divided by the previous matched level, less one. Applying the adjustment before computing returns prevents the unit change from becoming a spurious investment gain.

Checks cover positive input levels, conflicting equity observations, complete monthly coverage and the absence of an implausible liquid-fund return jump at the adjustment boundary. The source files are retained separately from transformed observations. Daily exchange-calendar completeness was not independently audited. The equity series is a total-return benchmark without an investable-product fee; the fund NAV reflects fund operating expenses. The latter is a market-valued liquidity proxy, not a guaranteed bank balance. Investor taxes and redemption frictions are excluded from the main specification.

3.2 Household profiles and observed benchmarks

The original household has monthly spending of Rs40,000, normal income of Rs60,000, initial investable wealth of Rs600,000 and an initial goal of Rs2,000,000. It provides continuity with the original experiment and is not labelled a representative household. Five additional profiles use a common spending scale of Rs27,984. This equals the HCES 2023–24 urban monthly per-capita consumption expenditure of Rs6,996, excluding imputation of specified free welfare items, multiplied by an assumed household size of four [11, Statement 1, p. 4]. This is a scale anchor rather than an observed essential-expense cash budget.

Three profiles set normal income to spending divided by one minus an assumed saving share of 5%, 10% or 20%. Two further profiles use the PLFS April–June 2024 mean urban regular-salaried individual earnings of Rs24,434 [12, Statement 10, p. 18]. One or two earnings benchmarks, each reduced by an assumed 10%, create the illustrative spendable incomes. The earner count and deduction are experimental assumptions; the deduction is not a statutory tax rate. Combining separate survey means does not estimate a joint household income, consumption and wealth distribution.

Table 1: Synthetic Household Starting Budgets

Profile	Spending Rs per month	Income Rs per month	Surplus Rs per month	Saving share %
Original illustration	40,000	60,000	20,000	33.3
Urban anchor with 5% saving	27,984	29,457	1,473	5.0
Urban anchor with 10% saving	27,984	31,093	3,109	10.0
Urban anchor with 20% saving	27,984	34,980	6,996	20.0
Two earnings benchmarks	27,984	43,981	15,997	36.4
One earnings benchmark	27,984	21,991	-5,993	-27.3

Sources: published scale anchors [11, 12] and stated assumptions. Values are rounded for display; calculations use unrounded inputs. Saving share equals one minus spending divided by income. The one-earner profile has an ongoing deficit before any interruption.

Initial wealth is fixed at 15 months of baseline spending and the initial goal at 50 months. Thus the added profiles begin with Rs419,760 of wealth and a Rs1,399,200 goal. All seven reserve policies are affordable at inception. These endowments and targets are assumptions, not survey estimates. The 2023–24 benchmarks provide the same synthetic starting scale in every historical return window; they do not reconstruct actual nominal budgets in the corresponding calendar years. Equal resources are maintained across reserve strategies within each profile and scenario, not across profiles.

4. Methods

4.1 Historical design and scenario structure

The model is a deterministic historical simulation, not a Monte Carlo draw of future returns. Each of the 85 consecutive starting months from January 2014 to January 2021 defines a 60-month return window. Adjacent windows share 59 months. Each scenario is evaluated at reserve sizes b in $\{0, 3, 6, 7, 8, 9, 12\}$, measured in months of normal spending. Forty-five scenario definitions yield $45 \times 7 \times 85 = 26,775$ paths. The paths are conditional policy experiments rather than independent household observations.

The original 27 definitions include a no-shock reference, interruptions of different duration and timing, partial income loss, lower normal income, temporary spending reductions, repeated interruptions, inflation and return-cost sensitivities, and proportional withdrawals. Eighteen further definitions apply selected interruptions to the five added profiles. Appendix A specifies this structure. The main comparison is full income loss in months 12 to 17. Monthly numbering is relative to each window, so the same interruption is placed against different observed market sequences. No unemployment hazard or joint distribution of income and market shocks is estimated.

Normal income, spending and the financial goal grow at an assumed 5% annually unless changed by a stated sensitivity. Income and spending step up at months 13, 25, 37 and 49. For normal spending E_0 and income Y_0 , their month- t values are $E_0(1 + g_E)^{\lfloor (t-1)/12 \rfloor}$ and $Y_0(1 + g_Y)^{\lfloor (t-1)/12 \rfloor}$. During a shock, income is multiplied by one minus the specified loss fraction, and spending by one minus the specified spending-cut fraction. Income resumes its normal path after the shock; there is no lasting earnings scar. The monthly-expense sensitivity uses exponent $(t-1)/12$ for spending only.

4.2 Monthly cash flows and asset allocation

Let C_t and Q_t denote closing liquid and equity assets, and let $r_{L,t}$ and $r_{E,t}$ be the observed monthly returns. Initial liquid assets equal bE_0 , while initial equity equals $W_0 - bE_0$. Returns are applied to opening assets before income and spending occur at month end. Write $C_t^* = C_{t-1}(1 + r_{L,t})$ and $Q_t^* = Q_{t-1}(1 + r_{E,t})$. The 1% annual equity-cost sensitivity additionally multiplies Q_t^* by $(1 - 0.01)^{(1/12)}$. The zero-cash-return sensitivity sets $r_{L,t}$ to zero. Both are stated sensitivity assumptions.

Let Y_t be income after the interruption and P_t planned spending after any temporary cut. The deficit is $D_t = \max(P_t - Y_t, 0)$ and the surplus is $S_t = \max(Y_t - P_t, 0)$. Under the cash-first rule, liquid withdrawals, equity sales and unmet planned spending are:

$$\begin{aligned} L_t &= \min(C_t^*, D_t); \\ X_t &= \min(Q_t^*, \max(D_t - L_t, 0)); \\ U_t &= \max(D_t - L_t - X_t, 0). \end{aligned}$$

The current reserve target is $B_t = bE_t$, using normal spending before any temporary cut. Remaining liquid assets are $A_t = C_t^* - L_t$. A positive surplus first replenishes the reserve by $F_t = \min(S_t, \max(B_t - A_t, 0))$. Any remaining liquid amount above target, $H_t = \max(A_t - B_t, 0)$, is transferred to equity. Closing balances are:

$$C_t = A_t + F_t - H_t;$$

$$Q_t = Q_t^* - X_t + S_t - F_t + H_t. \quad (1)$$

Equity is never sold merely to refill cash. This matters because replenishment competes with subsequent equity contributions rather than imposing a forced asset rebalance. Income, spending and the remaining surplus are accounted for once. In every month the balances satisfy:

$$C_t + Q_t = C_t^* + Q_t^* + Y_t - P_t + U_t. \quad (2)$$

Unmet planned spending represents consumption that is not funded. It is not borrowing and does not become a liability repaid in later months. A non-negative terminal balance therefore cannot establish that all earlier spending needs were met. The model also excludes informal support, credit, unemployment insurance and other compensating income unless represented by the stated surviving-income fraction.

For the proportional-withdrawal benchmark, the same opening allocation, returns, budgets and surplus-refill rule are retained. When total assets $V_t^* = C_t^* + Q_t^*$ are positive, the funded deficit $\min(D_t, V_t^*)$ is withdrawn from each asset in proportion to its value after returns. Liquid withdrawals are $\min(D_t, V_t^*)C_t^*/V_t^*$, with the remainder taken from equity. With no assets, the deficit is unmet. Later allocations differ endogenously because the withdrawal order differs. The comparison isolates that rule within this model rather than holding all future allocations constant.

4.3 Outcomes and goal definitions

Liquidation outcomes include cumulative nominal equity sales to fund spending, the share of windows with sales exceeding Rs0.01, and sales while the index is below its expanding observed peak. A deep drawdown is at least 20% below that peak. The peak is computed over the available matched history beginning in December 2013, not separately reset at each investment-window start. These are market-state measures: selling below a peak is not necessarily selling below the acquisition cost of the units, and realised tax-lot gains or losses are not computed.

Wealth outcomes include terminal total assets $W_{60} = C_{60} + Q_{60}$, mean and median wealth, and a lower-tail mean. For 85 ordered wealth outcomes, the lowest eight are weighted fully and the ninth at one-half weight; the total is divided by 8.5. This gives a finite-sample mean for the lowest 10%, not an estimated tail distribution for future markets. Monetary amounts are nominal rupees. Terminal wealth includes both initially invested resources and accumulated net contributions, so it is not a return rate.

The original goal test is $W_{60} \geq G_{60}$, where $G_{60} = G_0(1 + gG)^5$. The extension reports two reserve-adjusted tests: $W_{60} \geq G_{60} + bE_{61}$ for the strategy's own reserve, and $W_{60} \geq G_{60} + 6E_{61}$ for a common reserve. Here $E_{61} = E_0(1 + gE)^5$ incorporates the next annual spending increase. The principal extended outcome also requires cumulative unmet spending to be no more than Rs0.01:

$$\text{Protected success} = 1 \text{ when } W_{60} \geq G_{60} + 6E_{61} \text{ and } \sum U_t \leq \text{Rs}0.01. \quad (3)$$

Every policy faces the same terminal obligation in equation (3), including the policy that begins with no reserve. Six months is a reporting convention, not an estimated optimum. These are affordability tests permitting cost-free conversion between terminal assets; they do not require the retained reserve already to be entirely liquid. The own-reserve test is secondary because it imposes different end obligations on different strategies. The common-reserve extension was specified after inspecting the original results and

was not preregistered. The original goal result is retained, and sensitivity to initial goals of 30, 40 and 50 expense months is reported without changing return paths.

Reserve recovery is measured from the final shock month to the first month at or after it in which liquid assets reach the prevailing reserve target. Immediate restoration is coded as zero months. A path that does not restore by month 60 is right-censored and reported separately. The mean recovery duration includes restored paths only, and the restoration share accompanies it. No-shock paths have no post-shock recovery duration. For repeated interruptions, recovery is measured after the final interruption.

4.4 Reconciliation and interpretation

The implementation checks the monthly accounting identity, non-negative asset balances and consistency of nested goal definitions. The retained validation includes 28 zero-return fixtures, 141 selected paths checked against the workbook's monthly audit model, and exact reproduction of the packaged path and summary CSV files. The maximum monthly accounting residual is below Rs0.000001. A separate aggregation audit reconciled 23,695 comparisons across frozen results, saved workbook values, the earlier report and a worked example, with no failed comparison. This is internal numerical validation; it does not establish the empirical realism of the assumptions.

Comparisons are made within matched windows and resources. Shares are descriptive frequencies across the available windows. Overlap, repeated use of market observations and deliberately selected scenarios prevent interpretation as independent trials, estimated job-loss frequencies or future success probabilities. No independent-sample significance tests, causal estimates or utility rankings are reported. A zero-sale threshold means the smallest eligible value in the tested reserve grid; untested intermediate sizes may change its exact location.

5. Results

5.1 Liquidity protection and terminal wealth

Table 2 reports the main interruption for the original household. A six-month reserve reduces mean equity sales from Rs250,000 without a reserve to Rs5,777. Nevertheless, every window still includes a small sale at six months. Seven months is the smallest tested reserve that avoids sales in all 85 windows. All seven policies fund planned spending without unmet spending in this particular scenario, including zero reserve, which finances the interruption from equity. Avoiding liquidation and preventing unfunded consumption therefore produce different assessments.

Table 2: Reserve Size and Investment Outcomes After a Six Month Income Loss

Reserve months	Mean ending wealth Rs	Mean equity sales Rs	Windows with equity sales %	Lowest mean wealth Rs 10%
0	2,394,214	250,000	100.0	1,781,945
3	2,333,963	128,777	100.0	1,794,847
6	2,278,077	5,777	100.0	1,817,390
7	2,256,367	0	0.0	1,820,746
8	2,233,882	0	0.0	1,823,372


Reserve months	Mean ending wealth Rs	Mean equity sales Rs	Windows with equity sales %	Lowest 10% mean wealth Rs
9	2,211,398	0	0.0	1,825,998
12	2,143,961	0	0.0	1,833,878

Source: frozen simulation results, scenario 2. Original household; full income loss in months 12–17; cash-first withdrawals; 85 overlapping five-year windows per row. Sales fund planned spending. No policy has unmet spending in this scenario.

The six-month reserve starts at Rs240,000, but the interruption spans the first annual spending increase. Planned spending is Rs40,000 in month 12 and Rs42,000 in each of months 13 to 17, totalling Rs250,000. Liquid-fund returns offset part of this difference. The residual illustrates why a reserve labelled in initial expense months need not cover the same number of months after a budget increase. The result is conditional on the timing and inflation convention, not a general defect in six-month reserves.

Mean terminal wealth falls from Rs2,394,214 at zero reserve to Rs2,278,077 at six months, a difference of Rs116,137. In contrast, the lowest-10% mean rises from Rs1,781,945 to Rs1,817,390, an improvement of Rs35,445. Further increases in reserve size continue to reduce mean wealth while increasing this lower-tail measure in the reported scenario. Figure 1 displays the different directions of these outcomes. With no income interruption, the six-month policy lowers mean terminal wealth by Rs135,852 relative to zero reserve. The cost of liquidity is therefore present even when the reserve is not needed for a shock.

Figure 1: Liquidity Protection and Wealth Outcomes Under the Main Interruption



Source: the same scenario and observations as Table 2. One lakh equals Rs100,000. Mean and lower-tail wealth measure different objectives; neither curve is a welfare ranking.

At six months, 68.2% of windows contain a drawdown-related sale, with mean drawdown-related sales of Rs3,807. Only 3.5% contain sales during a drawdown of at least 20%. These categories locate withdrawals in the available index history. They do not establish realised capital losses, and their small amounts at six months should be read alongside the large reduction in total liquidation.

5.2 Goal affordability with a retained reserve

Table 3 shows why goal definitions affect the interpretation. The original goal of Rs2,000,000 grows to

Rs2,552,563 at month 60. Normal monthly spending at month 61 is Rs51,051, so the common reserve is Rs306,308 and the combined obligation is Rs2,858,871. These amounts are common to all reserve policies in this scenario. The reserve-protected goal-success share is consequently no higher than the original total-wealth goal share.

Table 3: Historical Goal Success Under Alternative Terminal Obligations

Reserve months	Total wealth goal %	Goal plus own reserve %	Goal plus common reserve and no unmet spending %
0	28.2	28.2	4.7
3	21.2	9.4	0.0
6	12.9	0.0	0.0
7	8.2	0.0	0.0
8	3.5	0.0	0.0
9	3.5	0.0	0.0
12	0.0	0.0	0.0

Source: scenario 2, 85 windows per row. The common reserve equals six months of month-61 normal spending. All paths here have no unmet spending, so the no-unmet condition does not cause the difference between the total-wealth and common-reserve columns. The common-reserve definition is a disclosed post-analysis extension.

For the six-month policy, total-wealth goal success is 12.9%, but no window meets the goal while retaining the common reserve. At zero reserve, the respective shares are 28.2% and 4.7%. Requiring a retained reserve changes the obligation rather than the historical returns. Because the initial wealth is fixed, larger reserve allocations also reduce initial equity exposure, influencing the later resources available for the goal. Liquidity protection should therefore not be described as improving goal affordability in this experiment.

Goal ambition has a separate effect. With the six-month policy and the same full interruption, the protected success shares for initial goals of 30, 40 and 50 expense months are 95.3%, 43.5% and 0%, respectively. These alternatives reuse the same asset paths. They demonstrate sensitivity to the stated goal rather than a reason to select the most favourable target retrospectively. In the two-earner illustration, the six-month policy has 49.4% total-wealth goal success and 8.2% success with the common reserve retained, showing that the size of the household surplus materially changes affordability.

5.3 Spending capacity and reserve recovery

Table 4 separates funding the interruption from restoring the reserve. For the original household, every window rebuilds the six-month reserve after the full interruption, taking 12.5 months on average from its final month. No path with a 5% or 10% initial saving share restores the reserve by month 60, although all planned spending is funded. With a 20% saving share, every path restores it, taking 24.4 months on average. The two-earner illustration restores in 11.0 months.

Table 4: Reserve Restoration After Full Income Loss by Household Profile

Profile	Mean ending wealth Rs	Unmet spending %	Restored by month 60 %	Mean recovery months
Original illustration	2,278,077	0.0	100.0	12.5
Urban anchor with 5% saving	605,982	0.0	0.0	Not restored
Urban anchor with 10% saving	727,075	0.0	0.0	Not restored
Urban anchor with 20% saving	1,025,819	0.0	100.0	24.4
Two earnings benchmarks	1,758,092	0.0	100.0	11.0
One earnings benchmark	25,840	60.0	0.0	Not restored

Source: scenarios 2, 28, 32, 36, 40 and 44. Six reserve months; full loss in months 12–17; 85 windows per row. Unmet spending is the share of windows with any unmet amount above Rs0.01. Recovery means include restored paths only. For an un-restored path, follow-up after month 17 is 43 months. Resources differ across profiles as defined in Table 1.

The low-saving profiles still make some equity sales, averaging Rs4,042, despite having no unmet spending. Their inability to rebuild does not mean that the interruption was unfunded, nor does funding it imply a full recovery. Annual increases in the reserve target and the small available surplus constrain restoration. The distinction matters when a subsequent disruption occurs before the liquid reserve has returned to its target.

Surviving income shortens recovery. Under a 50% loss for the same six months, the 10% saving profile restores in every window, taking 25.1 months on average. The corresponding means are 10.3 months at a 20% saving share and 3.0 months for the two-earner illustration. The 5% saving profile still does not restore within the horizon. A half-loss can illustrate one of two equal earners losing income, but the experiment does not estimate the incidence or dependence of earners' employment risks.

The one-earner profile is a boundary case with planned spending of Rs27,984 against income of Rs21,991 before a shock. It sells equity even in the no-shock scenario. With six reserve months and the full interruption, 60.0% of windows have unmet planned spending, and mean terminal wealth is Rs25,840. An emergency reserve delays or redistributes asset depletion under this recurring deficit; reserve size alone cannot make the normal budget sustainable.

5.4 Withdrawal order with identical initial resources

Table 5 compares cash-first and proportional withdrawals at a six-month reserve for full six-month losses starting in months 12, 30 and 48. Cash-first produces higher mean ending wealth in each comparison and higher wealth in 65.9% to 74.1% of paired windows. It also reduces spending-funded equity sales. These comparisons keep initial allocation and the surplus-refill rule fixed, so they address a separate choice from the size of the starting reserve.

Table 5: Cash First Minus Proportional Withdrawal Outcomes

Shock start month	Mean difference Rs wealth	Cash first higher wealth windows %	Mean equity sales avoided Rs
12	16,232	68.2	177,420
30	19,390	74.1	217,443
48	15,449	65.9	246,638

Source: paired scenarios 2/23, 4/24 and 6/25, six reserve months and 85 matched windows per pair. Positive wealth and sales-avoidance differences favour cash-first under these rules. A wealth win requires a difference above Rs0.01. Initial assets, returns and budgets match within each pair.

The cash-first rule does not win in every window. Its average advantage reflects the subsequent sequence of returns and allocations after withdrawals, including the diversion of new surpluses toward replenishment. These results cannot establish dominance over periodic rebalancing, borrowing-assisted policies or a utility-based strategy that responds to changing circumstances.

5.5 Scenario sensitivity and repeated interruptions

The smallest tested reserve that eliminates equity sales varies with the interruption and the ability to recover (Table 6). A 50% loss requires three months in the tested grid, while a 20% temporary spending reduction during the full loss requires six. Two full losses beginning in months 12 and 24 require twelve months; moving the second loss to month 30 lowers the threshold to seven. The intervening surplus therefore matters as well as the total number of lost-income months.

Table 6: Sensitivity of the Smallest Tested Reserve Avoiding All Equity Sales

Scenario variation	Reserve months
50% income loss for six months	3
20% spending cut during full loss	6
Full six month losses at months 12 and 24	12
Full six month losses at months 12 and 30	7
Monthly expense increases	7
Zero liquidity return	7
One percent annual equity cost	7
Eight percent annual expense growth	7
Full twelve month loss at month 12	None tested
Lower income and two twelve month losses	None tested

Source: original-household threshold results. Tested reserves: 0, 3, 6, 7, 8, 9, 12 months. Defaults: full six-month loss from month 12 and cash-first withdrawals.

Seven months remains the zero-sale threshold when expense increases are monthly, liquidity returns are set to zero, an annual 1% equity cost is imposed, or expense growth rises to 8%. This stability is specific to the liquidation criterion. The six-month policy's mean ending wealth falls to Rs2,188,455 with zero

liquidity return, Rs2,214,295 with the equity-cost assumption and Rs2,072,519 with 8% expense growth. Its original total-wealth goal-success shares are 4.7%, 4.7% and 0%, respectively. A stable reserve threshold does not imply stable investment-goal outcomes.

No tested reserve eliminates all equity sales during a full twelve-month interruption beginning in month 12. Nevertheless, those scenarios have no unmet spending because equity remains available. In the combined stress case, normal income is Rs45,000 and full twelve-month losses begin in months 12 and 30. No tested reserve avoids unmet spending in every window. At six reserve months, 96.5% of windows contain unmet spending. This severe scenario illustrates the limit of reallocating a fixed asset stock when the cumulative resource deficit is large.

6. Discussion

6.1 What an emergency reserve accomplishes

The results identify a trade-off rather than an investment improvement that holds across objectives. Maintaining liquid assets changes the timing and composition of withdrawals, reducing equity liquidation in the main interruption. It also leaves less of the initial wealth in equity and directs subsequent saving toward cash restoration. In the observed market sample, that combination lowers mean terminal wealth but improves the lowest-10% mean. Selecting a policy would require a stated preference over these outcomes and the costs of spending shortfalls. The model supplies the conditional outcomes needed for that discussion; it does not estimate those preferences.

The zero-reserve policy's higher mean wealth is also insufficient as a stand-alone recommendation. It relies on equity sales to finance the specified interruption. A household unwilling or unable to sell investments promptly may value liquid reserves beyond the measured terminal-wealth effect. Conversely, a sale below an index peak does not prove an economic loss relative to every alternative use of the money. Behavioural comfort, redemption timing and tax-lot losses would require additional evidence rather than inference from the drawdown indicator.

6.2 Goal obligations and recovery capacity

The common-reserve goal test prevents a household from counting the same resources as both spendable goal money and retained emergency protection. Applying a common terminal obligation also avoids giving the zero-reserve policy an easier goal solely because it starts without cash. The large change in reported goal success is therefore an accounting clarification with practical consequences. It does not mean that the model has generated new returns or discovered a statistically significant fall in success. Both the original and extended outcomes remain necessary to understand the comparison.

The household extensions show that savings capacity is a central conditioning variable. A reserve may successfully bridge a temporary interruption yet remain below target for the rest of the planning horizon. In that case, the household is less prepared for another interruption even though the original spending was funded. The repeated-shock scenarios reinforce this mechanism: a longer interval permits additional rebuilding and can lower the tested zero-sale threshold. Looking only at the duration of one loss misses the interaction among recovery time, normal surplus and shock spacing.

These findings suggest a structured planning use for the experiment. First specify normal spending, surviving income, accessible initial assets and the financial goal. Then examine spending coverage, liquidation, reserve restoration and goal resources as separate outcomes. The analysis is most relevant to asset-owning households that can choose the composition of existing financial wealth. The ongoing-deficit

profile demonstrates why a larger liquidity allocation cannot substitute for resolving a persistent mismatch between income and spending.

6.3 Relation to the reviewed literature

The findings are consistent with an established opportunity cost of holding emergency reserves [3], while adding Indian return histories and a common terminal reserve obligation to the comparison. The recovery analysis applies an existing distinction between crisis depletion and rebuilding [8] to financial-goal accumulation. The rural Indian studies [4, 5] remain important boundaries: their broader asset portfolios and coping mechanisms are outside the present two-asset experiment. The contribution is consequently the transparency and joint evaluation of outcomes within a specified application, rather than a new general theory of precautionary saving.

7. Limitations

First, the market sample is finite and contains one equity index and one selected liquid fund. It does not include all earlier Indian market episodes, alternative liquidity products or future return regimes. Eighty-five overlapping windows provide sequence variation but not 85 independent five-year histories. Selecting an available fund also limits product generality. There is no out-of-sample validation or resampled return uncertainty, and the reported shares should not be interpreted as calibrated forecasting probabilities.

Second, household profiles are synthetic. Published marginal spending and earnings benchmarks do not establish the joint distribution of income, wealth, household size, financial-market participation or unemployment risk. Initial wealth of 15 expense months assumes prior asset accumulation and excludes households without those resources. The expenditure anchor includes consumption beyond unavoidable cash bills. Fixed 5% growth and the optional 8% spending path are assumptions rather than reconstructed household inflation. No population weighting is applied.

Third, interruptions are imposed, not estimated. The design does not model job-search duration, permanent earnings damage, anticipatory saving, age, health shocks or dependence between employment loss and equity markets. The selected start months produce conditional stress tests. Full versus partial income loss is a transparent scenario choice, and the two-earner interpretation does not estimate correlated employment risks. The available reserve grid is coarse, so reported thresholds are minima among tested policies rather than exact continuous solutions.

Fourth, monthly sequencing assumes that returns accrue before income, spending and withdrawals at month end. Intramonth needs, settlement delays, taxes, exit charges, actual index-investment expenses and transaction restrictions are absent from the main model. A general equity-cost sensitivity addresses only part of this implementation gap. Credit, social insurance, family transfers and other buffers are omitted. Unmet spending is not assigned a utility penalty or carried forward as debt; it must therefore be assessed alongside wealth rather than folded into a single welfare score.

Finally, the retained-reserve goal is a post-analysis extension, and its six-month floor is a reporting choice. Recovery averages omit censored durations, making the accompanying restoration shares essential. The withdrawal benchmark does not span all investment policies, and the literature review is targeted. Internal reconciliation strengthens computational reproducibility but cannot correct these design limits. A broader claim would require richer household data, additional asset histories, explicit income-risk calibration and a stated welfare or decision criterion.

8. Conclusion

Emergency-reserve policies affect more than the ability to fund a single interruption. In the main historical simulation, a six-month reserve sharply reduces spending-funded equity sales and improves the lower tail of terminal wealth, while lowering average terminal wealth. Seven months is the smallest tested reserve eliminating sales for the original six-month loss beginning in month 12. That threshold changes with spending adjustment, surviving income and the spacing of repeated interruptions.

Two additional outcomes materially change the assessment. Retaining a common reserve at the goal date makes the investment obligation more demanding: the original household's six-month policy moves from 12.9% total-wealth goal success to no protected-goal success in the observed windows. Low saving capacity can also prevent reserve restoration even when all planned spending is funded. These distinctions caution against treating immediate liquidity coverage, goal affordability and financial recovery as interchangeable measures.

The study provides a reproducible applied comparison for synthetic asset-owning Indian households. It supports evaluating reserve size together with the normal surplus, surviving earnings, goal amount and recovery time. Its evidence is conditional on the specified cash flows and historical return sequences. A universally optimal reserve, population-wide outcome rate or future success probability is not established.

References

1. C.D. Carroll, "Buffer-stock saving and the life cycle/permanent income hypothesis", *The Quarterly Journal of Economics*, 1997, 112(1), 1–55. <https://doi.org/10.1162/003355397555109>
2. J.F. Cocco, F.J. Gomes, P.J. Maenhout, "Consumption and portfolio choice over the life cycle", *The Review of Financial Studies*, 2005, 18(2), 491–533. <https://doi.org/10.1093/rfs/hhi017>
3. J. Scott, D. Williams, J. Gilliam, J.P. Sybrowsky, "Is an all cash emergency fund strategy appropriate for all investors?", *Journal of Financial Planning*, 2013, 26(9), 56–62. <https://www.financialplanningassociation.org/article/all-cash-emergency-fund-strategy-appropriate-all-investors>
4. K.S. Imai, B. Malaeb, "Buffer stock savings by portfolio adjustment: Evidence from rural India", *Agricultural Economics*, 2015, 46(S1), 53–68. <https://doi.org/10.1111/agec.12198>
5. B. Malaeb, J. Golan, K. Imai, "Consumption smoothing and risk coping mechanisms: Evidence from rural India", *Review of Income and Wealth*, 2026, 72(2), e70060. <https://doi.org/10.1111/roiw.70060>
6. D. Farrell, F. Greig, C. Yu, "Weathering volatility 2.0: A monthly stress test to guide savings", *JPMorgan Chase Institute*, 2019. <https://www.jpmorganchase.com/institute/all-topics/household-financial-health/report-weathering-volatility-2-a-monthly-stress-test-to-guide-saving>
7. P. Ganong, P. Noel, "Consumer spending during unemployment: Positive and normative implications", *American Economic Review*, 2019, 109(7), 2383–2424. <https://doi.org/10.1257/aer.20170537>
8. A. Martin, M. Markhvida, S. Hallegatte, B. Walsh, "Socio-economic impacts of COVID-19 on household consumption and poverty", *Economics of Disasters and Climate Change*, 2020. <https://doi.org/10.1007/s41885-020-00070-3>
9. A. Gupta, A. Malani, B. Woda, "Explaining the income and consumption effects of COVID in India", *NBER Working Paper No. 28935*, 2021. <https://doi.org/10.3386/w28935>
10. J. Estrada, "The bucket approach for retirement: A suboptimal behavioral trick?", *The Journal of Investing*, 2019, 28(5), 54–68. <https://blog.iese.edu/jestrada/files/2019/09/BucketApproach.pdf>

11. Ministry of Statistics and Programme Implementation, “Fact sheet of Household Consumption Expenditure Survey 2023–24”, Government of India, 2024, Statement 1, p. 4. https://www.mospi.gov.in/sites/default/files/publication_reports/HCES%20FactSheet%202023-24.pdf
12. Ministry of Statistics and Programme Implementation, “Annual report Periodic Labour Force Survey July 2023–June 2024”, Government of India, 2024, Statement 10, p. 18. https://www.mospi.gov.in/sites/default/files/publication_reports/AnnualReport_PLFS2023-24L2.pdf
13. Nifty Indices, “Historical index data: Nifty 50 Total Return Index”, n.d., Observations December 2013–December 2025; collected 29 September 2026. <https://www.niftyindices.com/reports/historical-data>
14. Association of Mutual Funds in India, “Historical NAV records: HDFC Liquid Fund Direct Plan Growth, scheme 119091”, n.d., Observations December 2013–December 2025; collected 29 September 2026. <https://www.amfiindia.com/net-asset-value/nav-download>
15. HDFC Mutual Fund, “Addendum on change in face value of liquid schemes, effective 30 August 2015”, August 14, 2015. <https://files.hdfcfund.com/s3fs-public/addenda-reports/636%20-%20Addendum%20-%20Change%20in%20Face%20Value%20of%20Liquid%20Schemes%20-%20August%202014,%202015.pdf>
16. Source versions used. Reference [4] was reviewed through Kobe University Discussion Paper DP2014-14, revised 27 February 2015; reference [7] through NBER Working Paper 25417, January 2019; and reference [8] through the author preprint arXiv:2005.05945. The journal citations identify the associated publications without asserting that every detail of the reviewed working version is identical. References [1] and [2] provide abstract-level theoretical background.