International Reserves and Unspanned Import Risk

Milan Bumbulovic*

Working draft — August 2026

Abstract

International reserves are typically evaluated in monetary units even though one precautionary function of reserves is to preserve the capacity to finance imports during external stress. This paper evaluates conventional reserve portfolios directly in units of the external goods liabilities they may be required to finance. Using monthly data on major-currency government securities and gold, we measure liability-unit risk, solve constrained minimum-risk reserve portfolios, and translate the residual exposure into reserve-adequacy units. The choice of numeraire materially changes measured risk: a three-month U.S. Treasury bill has annualized volatility of about 1 percent in dollars but 27 percent in units of world energy, while conventional reserve assets become highly correlated when measured against the same goods liability. Yet reoptimization within the conventional reserve menu removes only about 0.5 percent of monthly world-energy variance. An economy beginning with three months of energy-import cover faces a fifth-percentile outcome of approximately 1.7 months after one year, with little variation across conventional portfolio allocations. The result survives alternative liability measures, portfolio restrictions, estimation methods, historical samples, and adverse energy-supply states. Expanding the opportunity set changes the result sharply: adding direct commodity-price exposure raises the hedgeable share substantially. The evidence therefore points to an instrument-set constraint rather than a portfolio-selection failure and distinguishes the monetary value of reserves from their purchasing power over specific external liabilities.

JEL classification: E58, F31, F33, F41, G11, Q43.

Keywords: international reserves; reserve adequacy; numeraire; import purchasing power; portfolio choice; oil supply shocks; commodity futures; strategic petroleum reserves.

1 Introduction

International reserves serve several purposes, including foreign-exchange intervention, external debt service, and protection against disruptions to external financing. For economies

*Independent researcher. Email: milan.bumbulovic@gmail.com. I thank Diego Känzig for the public oil supply news dataset. AI tools were used for research assistance, code, and data verification; analytical choices and text are my own. A replication package is available upon request. All errors are mine.

that depend heavily on imported energy, food, and other essential inputs, one important precautionary function of reserves is to preserve the capacity to finance imports when access to foreign exchange becomes more difficult. For that component of the reserve mandate, adequacy depends not only on the monetary value of the reserve stock but also on the quantity of external goods and obligations that the stock can finance.

Reserve-adequacy frameworks partly reflect this objective. Common indicators relate reserves to prospective imports, short-term external liabilities, and other potential demands on foreign exchange (International Monetary Fund, 2015; Independent Evaluation Office of the IMF, 2012). Reserve portfolios, however, are generally evaluated in monetary units. Expected returns, volatilities, and covariances are measured in dollars, euros, domestic currency, or a currency basket. There are good institutional reasons for doing so. Reserve assets are financial claims, and intervention, liquidity management, and debt service are themselves predominantly monetary functions. The question in this paper is therefore not whether monetary numeraires are inappropriate for reserve management. It is narrower: whether they fully characterize risk for the part of the reserve stock intended to preserve purchasing power over imports.

For that purpose, a monetary risk measure can omit an economically relevant relative-price exposure. The quantity of imports that a reserve portfolio can finance depends jointly on the value of the portfolio and the price of the goods being purchased. A portfolio that is stable in dollars can therefore be considerably less stable when measured in units of energy or other imports. This distinction is potentially most important during adverse external supply shocks. A supply disruption can raise the price of an essential import at the same time that the higher import bill increases the need for external financing. The reserve stock may therefore be called upon more intensively precisely when its purchasing power over the affected good has declined.

Sri Lanka's 2022 crisis provides an extreme illustration of the distinction. Foreign-exchange scarcity constrained the country's ability to finance fuel, medicines, and other imported necessities (International Monetary Fund, 2023b). Reserve portfolio design was not the cause of the crisis, and alternative portfolio weights within a conventional reserve mandate would not necessarily have prevented it. The narrower observation is that, once foreign exchange became scarce, the economic capacity provided by the remaining reserve stock depended on the imports that stock could purchase.

The exposure is also distinct from currency risk. Matching the currency composition of reserves to the currency in which imports are invoiced can reduce exchange-rate risk without hedging movements in the price of the underlying good. A dollar asset, for example, can match the payment currency of a dollar-invoiced oil import while leaving the holder fully exposed to changes in the dollar price of oil. This does not make currency matching ineffective: it addresses a different component of the liability. For dollar-invoiced trade, a dollar-heavy reserve portfolio may therefore be well aligned with intervention and payment needs while still leaving substantial exposure to movements in goods prices (Gopinath et al., 2020; Boz et al., 2022).

Existing research has recognized both sides of this problem. One literature relates reserve composition to trade, external liabilities, and import purchasing power; another studies commodity-linked instruments as hedges against sovereign commodity-price exposure (Ben-Bassat, 1980; Beck and Rahbari, 2011; Zhang et al., 2010a; Claessens and Varangis, 1991;

Daniel, 2001; Zhang et al., 2009). This paper builds on those approaches by asking a different question. Rather than beginning by adding a commodity allocation, we first hold the conventional reserve asset set fixed and ask how well that asset set can preserve purchasing power when its risk is evaluated directly in units of the liability. We then measure the portion of goods-price risk that remains after the best feasible portfolio has been selected.

The empirical analysis uses a long monthly panel of assets representative of the liquid core of reserve portfolios, including short- and longer-duration government securities across the major reserve currencies and gold. We evaluate these assets under conventional monetary numeraires and under goods-based numeraires. Energy is the principal case because it combines transparent internationally traded prices, large historical shocks, and limited short-run substitutability for many importing economies. Broader import-price and essentials-weighted baskets provide a more diversified measure of import purchasing power, while food is examined as an additional goods-price specification. The analysis does not attempt to construct a comparable quantitative medicine-price liability, for which internationally consistent long-run price data and procurement structures are substantially less tractable.

The empirical exercise proceeds in four stages. First, we measure how the volatility and covariance structure of conventional reserve assets changes when returns are expressed in monetary units and in units of imports. Second, we solve the reserve portfolio problem under standard portfolio constraints and measure the minimum purchasing-power risk that the conventional asset set can attain. We refer to this residual risk as the **hedging deficit**, and report the corresponding share of liability variance that the feasible reserve portfolio can hedge. Third, we translate these portfolio results into reserve-adequacy units by estimating the distribution of import cover, and separately calculate the certainty-equivalent cost of selecting portfolios under one numeraire when performance is evaluated under another. The latter exercise is treated as secondary because it depends on estimates of expected returns. Fourth, we examine whether the results change in adverse energy-supply states and in ex-ante event exercises, before expanding the permissible asset set to instruments with more direct exposure to energy prices.

The main results are threefold. First, the choice of measuring unit materially changes the measured risk of conventional reserve assets. A three-month U.S. Treasury bill has annualized volatility of roughly 1 percent in dollars but approximately 27 percent when its return is measured in units of world energy. More generally, assets that provide meaningful diversification in monetary units become highly correlated when measured against the same volatile goods liability.

Second, changing portfolio weights within the conventional reserve asset set does little to eliminate this exposure. At the monthly horizon, the minimum-risk portfolio removes only about one-half of one percent of the variance associated with the world-energy liability. The result is not a failure of the optimization procedure: when the same method is applied to a monetary liability whose components are represented in the feasible asset set, it is able to replicate a substantial share of that liability. The difference is therefore primarily one of available instruments rather than portfolio selection. The conclusion is also robust to conditioning on adverse supply states: the conventional menu continues to provide little insurance against the goods-price exposure in the states in which such insurance would be most relevant.

Third, the residual risk is economically meaningful when expressed in the same units used

to assess reserve adequacy. In the baseline exercise, an economy beginning with three months of energy-import cover faces a fifth-percentile outcome of approximately 1.7 months after one year, with relatively little variation across portfolios optimized under different numeraires. In contrast, expanding the opportunity set to include instruments whose payoffs are directly related to energy prices produces substantially larger reductions in purchasing-power risk, although the effectiveness of financial hedges remains limited by basis risk, roll costs, and reserve-management constraints.

These results should not be interpreted as evidence that central banks have selected the wrong reserve portfolios or the wrong reporting currencies. Conventional reserve assets are chosen under demanding requirements for liquidity, safety, market depth, and availability for intervention and external payments. Monetary benchmarks are well suited to several of those functions. The results instead identify a boundary of what those assets can accomplish. If part of the reserve stock is viewed as insurance against interruptions to essential imports, a material component of the relevant relative-price risk may lie outside the span of the conventional reserve portfolio. In that case, changing the portfolio benchmark may improve the measurement of the exposure without creating an asset capable of hedging it.

The final part of the paper examines that distinction directly. We progressively introduce financial instruments with greater exposure to energy prices and compare them with physical strategic reserves held elsewhere on the sovereign balance sheet. The purpose is not to argue that central banks should mechanically replace conventional reserve assets with commodities or derivatives, nor that physical stockpiles and foreign-exchange reserves are institutionally interchangeable. Rather, the exercise asks which risks can be managed within the conventional reserve portfolio and which require instruments elsewhere in the sovereign's broader external-insurance framework.

The contribution of the paper is therefore primarily one of measurement and portfolio feasibility. Asset risk depends on the unit in which performance is evaluated. For the import-insurance component of reserves, evaluating conventional reserve assets only in monetary units can omit a material relative-price exposure. We quantify that exposure, measure how much of it can be reduced within a reserve-eligible asset set, and identify how the answer changes when the opportunity set is expanded. This provides a way to evaluate the import-insurance function of reserves without requiring that the other objectives of reserve management be expressed in the same unit.

2 Reserve Management, Import Risk, and Related Literature

International reserves serve multiple purposes. They support foreign-exchange intervention, external debt service, liquidity provision, and precautionary protection against disruptions to external financing. The importance of the precautionary motive is well established in the literature, particularly for emerging and developing economies (Aizenman and Lee, 2007; Aizenman and Riera-Crichton, 2008; Aizenman et al., 2012). Reserve-adequacy frameworks similarly relate the reserve stock to potential demands on foreign exchange. The traditional import-cover measure compares reserves with prospective imports, the Guidotti-Greenspan

rule focuses on short-term external debt, and the IMF's reserve-adequacy framework incorporates several potential sources of external financing pressure (International Monetary Fund, 2015; Independent Evaluation Office of the IMF, 2012).

These measures do not imply that import purchasing power is the appropriate objective for the entire reserve portfolio. Intervention and debt-service needs are naturally defined in monetary terms, and different reserve tranches can serve different purposes. They do establish a narrower point relevant to this paper: for the precautionary component of reserves intended to support external purchases during periods of stress, the value of the assets is relevant partly in relation to the external liabilities they may need to finance.

Reserve-management practice already reflects this dependence of the relevant benchmark on the purpose of the reserves. In the BIS survey of reserve-management practices, central banks reported using domestic currency, individual foreign currencies, and currency baskets as portfolio numeraires, with some institutions applying different numeraires to different tranches (Borio et al., 2008). McCauley and Chan make the same point explicitly: intervention needs imply a benchmark related to intervention currencies, debt-service needs imply a benchmark related to the currency composition of liabilities, and protection of foreign purchasing power can imply an import-based benchmark (McCauley and Chan, 2014). Lu and Wang similarly incorporate import payments and short-term external debt into the determination of the liquidity tranche, with its currency composition reflecting the currencies in which those obligations are denominated (Lu and Wang, 2019).

The distinction examined here arises after this currency matching has taken place. Matching a dollar-denominated import liability with dollar assets reduces exchange-rate risk, but it does not stabilize the dollar price of the imported good. The relevant exposure is therefore not evidence that existing reserve-management practice ignores liabilities. Rather, it is a residual component of the liability that remains after its monetary denomination has been taken into account.

Historical energy shocks provide a simple illustration. During the 1973–74 oil shock, benchmark crude prices increased severalfold within a short period. A dollar asset that retained its nominal value consequently lost a large share of its purchasing power over oil even though the asset itself had experienced no comparable loss (Adelman, 1990; Hamilton, 2013). The 1979–80 episode shows why the source of the price increase also matters. The increase in oil prices reflected a combination of supply disruption, precautionary demand, inventory behavior, and broader demand conditions (Barsky and Kilian, 2002; Kilian, 2009). The covariance between reserve assets and energy prices need not be the same across these states, which motivates the state-dependent analysis later in the paper.

More recent episodes illustrate the same distinction under different financing conditions. In Sri Lanka in 2022, very low usable foreign-exchange reserves eventually constrained the financing of fuel and other essential imports (International Monetary Fund, 2023b). This does not imply that reserve composition caused the crisis or that alternative portfolio weights would have prevented it. It shows that, once external financing became binding, the purchasing power of the remaining reserve stock over essential imports became economically relevant. The euro area experienced a large deterioration in its energy import bill in the same year without a comparable shortage of foreign exchange. Higher energy prices instead appeared primarily through the terms of trade, the current account, and domestic income (European Central Bank, 2023). The exposure to the import-price shock was present in both

cases; their ability to finance it was very different. Figure 1 traces the same distinction over the full sample: a reserve stock earning the Treasury-bill return, valued in dollars and in units of world energy, together with realized import-cover paths for Sri Lanka and India. A brief account of these and related episodes is provided in Appendix A.

2.1 Related literature

The paper builds on several strands of work that have moved reserve management toward a liability- or purchasing-power-based objective.

A first strand considers the unit in which reserve performance is evaluated. Zhang Bin, Wang Xun and Hua Xiuping (2010) calculate returns on China's foreign-exchange reserves in dollars and construct corresponding effective-return measures using both a currency basket and a basket of goods and services. They show directly that the measured performance of the reserve stock depends on the unit in which it is evaluated (Zhang et al., 2010a). Their analysis establishes an important measurement precedent for the present paper. It does not, however, study the resulting portfolio choice across a broader reserve asset set.

A related reserve-composition literature connects optimal currency holdings to the structure of trade and external payments. Ben-Bassat and subsequent work link desired reserve currency shares to trade and liability exposures (Ben-Bassat, 1980). Beck and Rahbari come particularly close to the exercise considered here. Their baseline analysis studies optimal dollar and euro reserve shares, and a robustness exercise replaces consumer-price deflation with country-level import-price indices. The choice of deflator materially affects the optimal allocation for some countries, including sizeable changes in the dollar share for Türkiye and Brazil (Beck and Rahbari, 2011, 2008, §5.2.6, Table 6). Vaccaro-Grange similarly studies currency weights as a means of reducing the exchange-rate component of imported-inflation risk (Vaccaro-Grange, 2026).

This work establishes that trade exposures and import prices can affect the appropriate reserve benchmark. The present paper extends that logic along two dimensions. First, it applies goods-based measurement to a broader menu of conventional reserve assets, including duration exposure, foreign government bonds, and gold rather than only currency shares. Second, it asks not only which feasible portfolio has the lowest risk in that unit, but also how much of the goods-price exposure remains after that portfolio has been selected.

A second strand begins from the instrument side. The sovereign commodity-risk literature studies derivatives and other commodity-linked contracts as tools for managing public-sector exposure to commodity-price fluctuations (Claessens and Varangis, 1991; Daniel, 2001). This literature establishes that direct commodity exposure can be useful when the underlying fiscal or external liability is itself commodity-sensitive.

A closely related Chinese literature places commodity-linked instruments explicitly within the foreign-reserve allocation problem. Zhang Hailiang, Wu Chongfeng and Zou Ping (2009) propose allocating a portion of China's reserves to physical commodity assets in relation to domestic industrial demand (Zhang et al., 2009). Subsequent work studies commodity futures and commodity-related equities as additions to the reserve portfolio (Zhang et al., 2010b; Zhang and Chai, 2010). Zhang Hailiang (2011) develops an allocation intended jointly to preserve reserve returns and reduce fluctuations in the cost of imports, subject to reserve-management constraints (Zhang, 2011). Yin and Han (2014) continue this line of work on

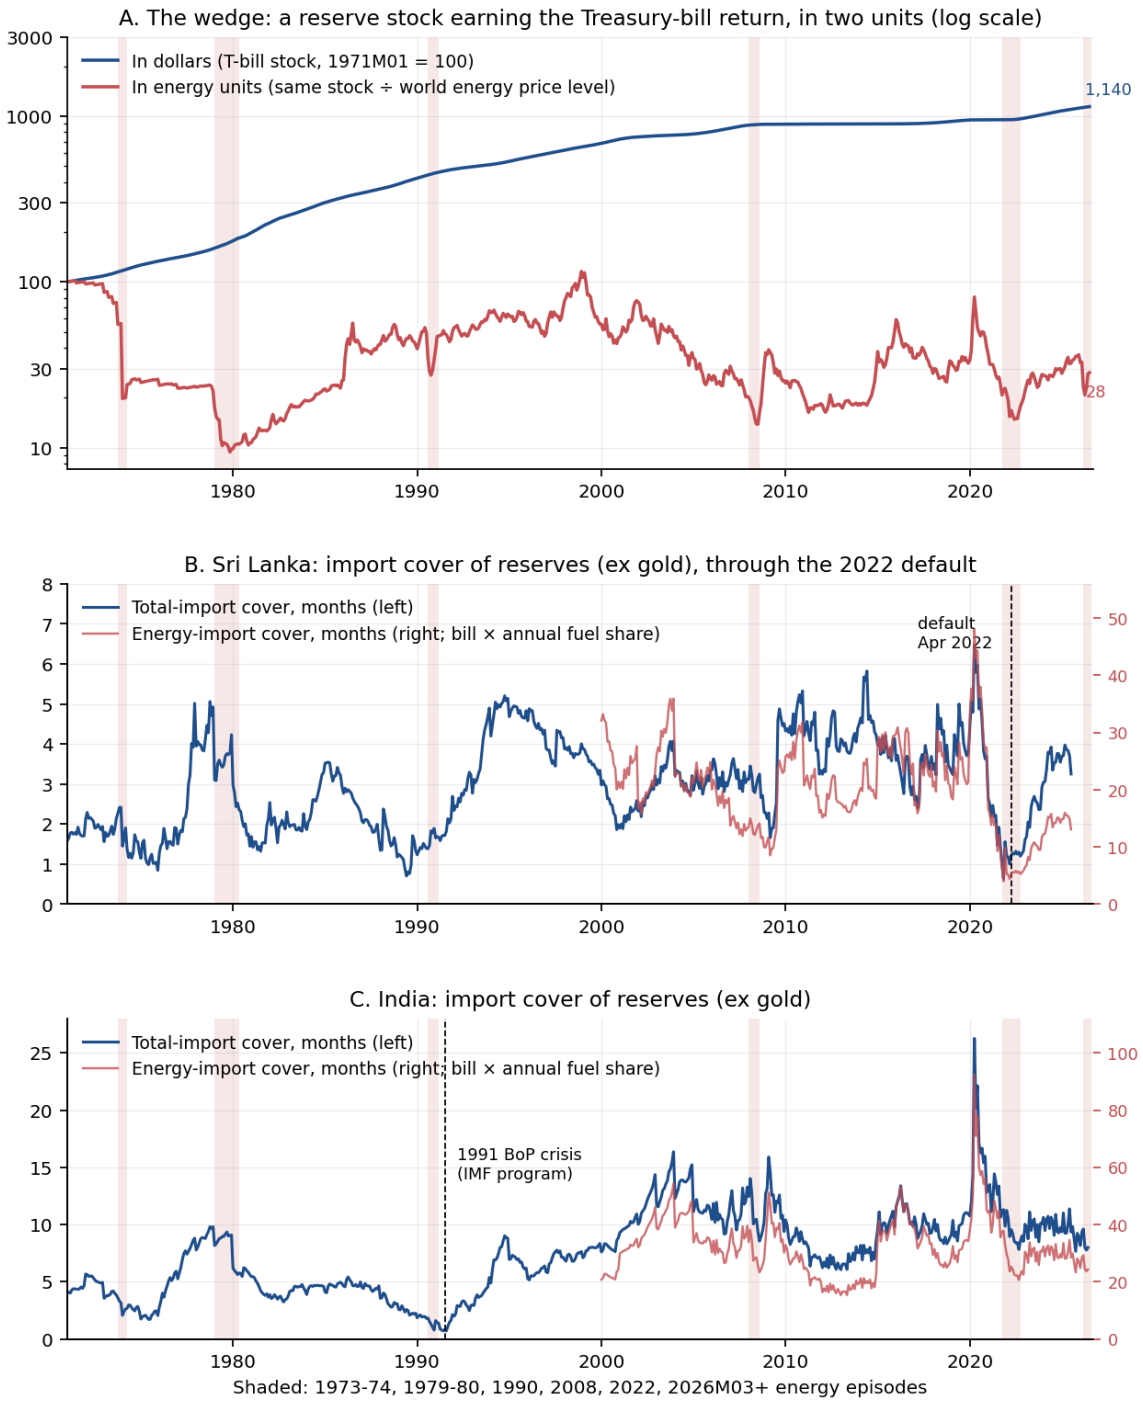


Figure 1: The wedge and two realized cases. Panel A: a reserve stock earning the Treasury-bill return, indexed to 100 in January 1971, valued in dollars and in units of world energy (log scale); the dollar line ends at 1,140, the energy line at 28, with a trough of 9.4 in November 1979. Panels B and C: realized import cover of reserves (excluding gold) for Sri Lanka through its 2022 default and for India, in months of total imports (left scale) and of energy imports (right scale). Shaded bands mark the 1973–74, 1979–80, 1990, 2008, 2022 and 2026 energy episodes. Sources: World Bank Commodity Markets (World Bank, 2026) for the world energy index; IMF international-liquidity series (International Monetary Fund, 2026b) for reserves and imports; construction in Appendix B and the replication files.

strategic commodity allocation (Yin and Han, 2014).

This literature identifies the instrument class that is most directly related to the exposure studied here. Our contribution is therefore not the proposition that commodities or commodity-linked claims can hedge commodity-price risk. Instead, we first restrict the portfolio to the conventional reserve asset set and quantify how much goods-price risk remains within that constraint. We then expand the opportunity set to commodity-linked financial instruments and compare the resulting improvement with the baseline. This sequencing allows the value of the additional instruments to be measured relative to the risk that conventional reserve assets leave unhedged.

Taken together, the existing literature approaches the problem from two directions. One changes the benchmark or objective so that reserve performance is evaluated relative to trade, imports, or broader purchasing power. The other adds instruments whose payoffs are directly related to commodity prices. The analysis below connects these two approaches: it changes the unit in which the conventional reserve portfolio is evaluated, measures the residual exposure, and then examines how that exposure changes when the asset set is expanded.

2.2 The conventional reserve asset set

The distinction between these two approaches matters because actual reserve portfolios remain concentrated in liquid monetary assets. Reserve mandates vary across institutions, and some central banks are permitted to hold a considerably broader range of securities. Actual holdings are more concentrated. In the 2023 World Bank RAMP survey, government bonds, bank deposits, money-market instruments, and supranational securities together accounted for approximately 78 percent of reported reserve portfolios, while corporate bonds and developed-market equities represented much smaller shares (World Bank, 2023). Reserve holdings are also concentrated by currency, with the dollar and euro continuing to account for the majority of reported foreign-exchange reserves (International Monetary Fund, 2026a). Gold is material for many reserve holders but remains a distinct component of the conventional reserve portfolio.

Reserve managers also commonly divide reserves into liquidity and investment tranches (Borio et al., 2008; World Bank, 2020, 2023). These institutional features matter for the empirical exercise because the relevant comparison is not with an unconstrained global investment portfolio. A central bank cannot generally respond to an identified import-price exposure by selecting any asset with a favorable historical covariance. Reserve assets must also satisfy requirements relating to liquidity, credit quality, market depth, operational capacity, and availability during periods of stress.

We therefore take the conventional liquid reserve portfolio as the initial opportunity set. The purpose is not to treat this set as universal across central banks, but to approximate the common core of assets that reserve managers can plausibly hold at scale. The specific instruments and portfolio constraints used in the empirical analysis are described in Section 4. Commodity futures, commodity-related financial assets, and physical strategic stocks are excluded from the baseline and introduced separately as extensions.

2.3 The relevant import liability

The final issue is the definition of the liability against which the portfolio is evaluated. Total imports need not provide the most relevant measure during a balance-of-payments or external-financing stress. When foreign exchange becomes scarce, discretionary imports can contract while authorities continue to prioritize goods that are difficult to postpone or substitute. The experiences of Sri Lanka and Pakistan provide examples of explicit prioritization of essential imports under foreign-exchange constraints (International Monetary Fund, 2023b,a; Khan and Knight, 1988).

The appropriate stress-period basket is therefore country-specific. The paper does not attempt to identify a universal set of essential imports. Instead, it uses energy as the baseline liability and then examines broader import-price measures. Energy is useful empirically because internationally traded prices are observable over a long sample, price shocks are large, and short-run substitution can be limited for importing economies. Energy prices can also affect other import costs through transportation, industrial production, and agricultural inputs (International Energy Agency, 2021; Levi and Molnar, 2022). Food is examined as an additional goods-price measure, while broader essentials-weighted and import-price baskets provide less concentrated measures of external purchasing power. Medicines are relevant to the underlying motivation but are not modeled as a separate quantitative liability because internationally comparable long-run price series and procurement structures are substantially more heterogeneous.

The role of energy in the analysis should therefore be interpreted narrowly. It is not assumed to represent the entire objective of reserve management or even the entire import-insurance objective. It provides a relatively transparent stress liability against which the consequences of changing the portfolio's unit of measurement can be evaluated. The broader import baskets indicate how the results change as the liability becomes more diversified.

The discussion above establishes three elements of the empirical problem. Reserve management has recognized liability-sensitive objectives and benchmarks; conventional reserve portfolios nevertheless remain concentrated in liquid monetary claims; and an important component of external purchasing-power risk can arise from movements in the price of the goods those claims may need to finance. The next section formalizes the portfolio problem when asset returns are evaluated relative to that liability.

3 Framework

The empirical exercise evaluates reserve assets in units of the liability they may be required to finance. Let $R^{\$}_{i,t}$ denote the dollar return on asset i, and let $R^{\$}_{N,t}$ denote the change in the dollar price of a measuring unit N. The return on the asset measured in units of N is

$$R^{N}_{i,t} = \frac{1 + R^{\$}_{i,t}}{1 + R^{\$}_{N,t}} - 1. \tag{1}$$

This transformation is used throughout the paper. When N is an import price or goods basket, $R^{N}_{i,t}$ measures the change in the quantity of that basket that the asset can purchase.

The empirical analysis applies the transformation to monetary benchmarks, including the dollar and SDR, and to energy, food, and broader import-price measures.

Equation (1) also makes clear why changing the measuring unit can change the risk characteristics of the portfolio. An asset with low volatility in dollars need not have low volatility relative to an import liability. Its risk in liability units depends on both its own return and its co-movement with the price of the liability. If the liability price is volatile and the asset does not tend to rise with it, much of that liability-price variation remains in the transformed return.

We apply this transformation to each asset in the feasible reserve set. Let R_t^N denote the vector of transformed asset returns and let

$$\Sigma^N = \mathrm{Cov}(R_t^N)$$

denote their covariance matrix. For measuring unit N, the minimum-risk reserve portfolio is

$$w_N^* = \arg\min_{w \in \mathcal{W}} w' \Sigma^N w, \tag{2}$$

where $\mathcal{W}$ denotes the feasible reserve portfolio. The specific assets and portfolio constraints defining $\mathcal{W}$ are described in the next section.

Equation (2) is a standard constrained minimum-variance problem. This is deliberate. The purpose of the exercise is not to introduce a new portfolio criterion, but to isolate the effect of changing the unit in which reserve risk is evaluated and the assets available to manage that risk. When N represents an import liability, the solution identifies the feasible reserve portfolio with the lowest volatility of purchasing power over that liability.

The main object of interest is the risk that remains after this optimization. We define the **hedging deficit** as

$$HD(N) = \sqrt{\min_{w \in \mathcal{W}} w' \Sigma^N w}. \tag{3}$$

The hedging deficit is therefore the minimum volatility of reserve purchasing power that can be achieved with the feasible asset set. A positive value does not by itself indicate inefficient reserve management. It can instead arise because the assets available to the reserve manager do not provide sufficient exposure to movements in the liability price.

For comparisons across liabilities, we also report a normalized measure. Let $\sigma^2_{0,N}$ denote the variance of the liability-unit return on a zero-return dollar position. The **hedgeable share** is

$$HS(N) = 1 - \frac{HD(N)^2}{\sigma^2_{0,N}}. \tag{4}$$

This statistic measures the proportion of liability-unit variance eliminated by the minimum-risk feasible portfolio relative to the dollar benchmark. A value of one indicates that the liability can be fully hedged within the feasible set. A value near zero indicates that changing portfolio weights does little to reduce the exposure. Under portfolio constraints, the measure can be negative if every feasible portfolio is more volatile in liability units than the benchmark.

The same relative-price calculation has a direct interpretation in reserve-adequacy terms. Let C_t denote the number of months of imports that the reserve stock can finance. Holding the quantity of imports fixed over the measurement horizon,

$$\frac{C_{t+H}}{C_t} = \prod_{\tau=1}^{H} \frac{1 + R^{\$}_{p,t+\tau}}{1 + R^{\$}_{N,t+\tau}}, \tag{5}$$

where $R^{\$}_{p,t}$ is the dollar return on the reserve portfolio. The evolution of import cover is therefore governed by the same ratio used to measure asset returns in liability units.

To express the resulting risk in a familiar adequacy measure, we report **Import-Months-at-Risk (IMaR)**. For the twelve-month horizon, let $G_{t,12}(w) = C_{t+12}/C_t$. Using three months of initial import cover as a common normalization,

$$IMaR_{5\%}(w) = 3\, Q_{0.05}(G_{t,12}(w))\,. \tag{6}$$

IMaR is the fifth-percentile number of months of import cover remaining after one year for a reserve position beginning at three months of cover. The three-month starting value is a normalization rather than an assumption about reserve adequacy; results for other initial levels scale proportionally.

The covariance between reserve assets and the liability may also depend on the source of the price movement. An increase in energy prices associated with stronger global demand need not generate the same asset-price response as an increase associated with an adverse supply shock. We therefore re-estimate the covariance structure and the portfolio problem conditional on the shock states defined in the empirical section. No particular state classification is imposed by the framework itself.

Welfare comparison. The headline analysis is based on risk rather than expected returns. As a secondary exercise, we also measure the cost of selecting a portfolio under one measuring unit when performance is ultimately evaluated under another. For risk-aversion parameter γ, define the certainty equivalent in unit N as

$$CE_N(w;\gamma) = w'\mu^N - \frac{\gamma}{2} w'\Sigma^N w, \tag{7}$$

where μ^N is the expected return vector in that unit. If w_N^{CE} maximizes this quantity under measuring unit N, cross-numeraire regret from choosing the portfolio appropriate to unit j when welfare is evaluated in unit k is

$$L_\gamma(j \to k) = CE_k(w_k^{CE};\gamma) - CE_k(w_j^{CE};\gamma). \tag{8}$$

We treat this calculation as secondary because it depends on expected-return estimates, which are less precisely estimated than the covariance structure. The estimation of expected returns and the sensitivity analysis used for this exercise are described in the next section.

The framework therefore separates three elements of the reserve problem: the liability against which purchasing power is evaluated, the feasible asset set, and the covariance between the two. The empirical analysis estimates these objects for conventional reserve portfolios and examines how the results vary across liabilities, horizons, and states.

4 Data and Empirical Strategy

The analysis uses monthly data from January 1971 through December 2025. Observations from 2026 are reserved for the out-of-sample exercises reported later in the paper. The dataset combines returns on assets representative of conventional reserve portfolios with monetary and goods-price measures used as alternative units of evaluation. Detailed definitions, sources, return construction, and series-specific adjustments are reported in Appendix B.

4.1 Assets and liability measures

The baseline reserve menu consists of U.S. Treasury bills, intermediate- and long-duration U.S. Treasuries, German, Japanese, and U.K. government bonds, and gold. The SDR is included as a diversified monetary benchmark rather than as an additional risky asset in the baseline portfolio. Government-bond total returns are constructed from official or fitted government yield curves. Foreign government bonds are left unhedged in the baseline specification because currency composition is itself part of the reserve-allocation decision. Currency-hedged variants are reported as robustness exercises.

The principal goods liability is the World Bank world energy price index. Energy provides the baseline case because it has a long monthly history, internationally observable prices, and repeated large price movements. It also represents an important component of essential imports for many energy-importing economies.

We consider several alternative liability measures. The World Bank food price index provides a second commodity-price specification. A broader import-price measure combines energy, food, and manufactured-goods prices, with weights intended to represent the composition of imports relevant during external stress. We also report a conventional trade-weighted version of this basket. Country-specific exercises use national or country-weighted measures where suitable data are available, including Türkiye's published import unit-value index. The detailed construction of each series and the alternative weighting schemes are described in Appendix B.

The empirical design therefore compares two monetary measuring units—the dollar and SDR—with several goods-based measures. The goods indices are not assumed to represent investable reserve assets. They define the liabilities against which the purchasing power of the reserve portfolio is evaluated.

4.2 Portfolio constraints and estimation

The baseline portfolio is long-only and fully invested. A minimum of 30 percent of the portfolio is allocated to U.S. Treasury bills to approximate the liquidity component of reserve portfolios documented in Section 2. The remaining allocation is chosen across the other assets in the conventional reserve menu. The baseline does not impose issuer caps; a specification limiting exposure to individual non-U.S. issuers is reported as a robustness exercise.

These restrictions are intended to approximate a conventional liquid reserve portfolio rather than the broadest portfolio that any central bank is legally permitted to hold. Commodity futures, commodity-linked securities, and physical strategic reserves are therefore

excluded from the baseline opportunity set. They are introduced separately when the asset menu is expanded later in the paper.

For each measuring unit, the return transformation defined in Section 3 is applied to all assets before the covariance matrix is estimated. Headline covariance estimates use Ledoit-Wolf shrinkage, with conventional sample covariance estimates reported for comparison (Ledoit and Wolf, 2004). The minimum-risk portfolio, hedging deficit, and hedgeable share are then calculated from the same transformed return panel.

Uncertainty is assessed using a moving-block bootstrap of the joint monthly return panel. The baseline implementation uses 24-month blocks, 1,000 replications, and 90 percent confidence intervals. Alternative block lengths are reported as robustness checks. Resampling joint return histories preserves the observed contemporaneous dependence between reserve assets and the liability price.

The horizon analysis considers one-, twelve-, twenty-four-, and sixty-month holding periods. Multi-period returns are compounded before being transformed into the relevant liability unit. Because observations overlap at horizons longer than one month, the bootstrap procedure is adjusted accordingly. A three-month specification is used as a diagnostic for differences in timing between end-of-month financial prices and goods-price series reported as monthly averages. Full implementation details are given in Appendix B.

Expected returns enter only the secondary welfare comparison defined in Section 3. We report that exercise under two mean-return specifications: sample means partially shrunk toward a common mean and a zero-excess-return specification relative to Treasury bills. The covariance-based results do not depend on these assumptions.

4.3 Identification of energy-price states

The state-dependent analysis asks whether the covariance between reserve assets and energy prices differs according to the source of the energy-price movement. The baseline identification uses the oil-supply-news shock of Känzig (2021), which measures revisions to expected future oil supply around OPEC-related news.

The headline conditional sample begins in July 1983, when the identifying series is available without model-based extrapolation. We standardize the published shock series and define an adverse supply state as a month in which the price-raising supply shock is positive and exceeds one-half of a standard deviation. A one-standard-deviation threshold provides a stricter alternative.

All remaining observations form the complement to the baseline supply state. Within that complement, the structural decomposition of Baumeister and Hamilton is used to distinguish demand-related oil-price movements where that distinction is required (Baumeister and Hamilton, 2019). The purpose of these classifications is not to impose a particular structural interpretation on the unconditional portfolio results. They allow the covariance matrices defined in Section 3 to be re-estimated in economically different energy-price states.

Several alternative classifications are used to assess sensitivity. These include an announcement-only implementation of the Känzig identification, alternative thresholds, a specification based directly on the structural oil-shock taxonomy, and an alternative implementation following the approach of Ready (2018). Pre-1983 values of the Känzig shock require extrapolation

and are therefore excluded from the headline conditional estimates. Details of the alternative classifications and their correspondence across series are reported in Appendix B.

4.4 Estimation samples

Different empirical exercises require different common samples because the underlying series begin at different dates. We use the longest available observations for descriptive asset statistics when doing so does not require direct portfolio comparisons. Portfolio optimization and the baseline hedging statistics use the common period in which all assets in the conventional reserve menu are available, beginning in October 1974, the first month in which returns are available for every asset in the conventional menu. The conditional analysis uses July 1983 through December 2025. Exercises that add instruments with shorter histories use the common sample available for those instruments.

The principal sample conventions are summarized in Table 1.

Table 1: Principal sample conventions.

Exercise	Sample convention	Main horizon(s)
Asset risk and pairwise relationships	Longest available common series for each comparison	1 month
Baseline portfolio, hedging deficit, hedgeable share	Full conventional-menu sample, October 1974–December 2025	1, 12, 24, 60 months
State-dependent analysis	July 1983–December 2025	Primarily 1 month, with horizon extensions
Import-Months-at-Risk	Historical joint return paths	12 months
Expanded financial asset menu	Common sample available for added instruments	1 and 12 months
Event/out-of-sample exercises	Formation periods specified for each exercise	Event-specific

The use of different samples is therefore determined by data availability and identification rather than by the realized results. Appendix B reports the exact start and end date used for each table and figure.

The empirical design is intentionally uniform across measuring units. The same asset returns, portfolio restrictions, covariance estimator, and optimization procedure are applied after changing only the liability unit against which the portfolio is evaluated. Differences in the resulting portfolios and residual risk can therefore be attributed to the change in the relevant liability and its covariance with the feasible reserve asset set.

5 Main Results

This section presents the main empirical results for the conventional reserve asset set defined in Section 4. The analysis proceeds in four steps. Section 5.1 documents how the measured risk of reserve assets changes when returns are evaluated relative to alternative liabilities.

Section 5.2 asks whether portfolio reallocation within the conventional reserve menu can materially reduce the resulting exposure. Section 5.3 translates the residual risk into months of import cover. Section 5.4 examines whether the conclusions change in adverse energy-supply states.

5.1 Risk Reclassification Across Numeraires

Table 2 reports annualized monthly volatility for representative reserve assets under alternative measuring units. The first two columns use monetary numeraires; the remaining columns evaluate the same assets relative to goods-price liabilities. All returns are transformed using Equation (1).

Table 2: Asset volatility by measuring unit. Annualized percent, monthly horizon. The table reports the baseline zero-coupon convention for U.S. Treasury securities; alternative return constructions and the full asset-by-numeraire grid are reported in Appendix C.

	USD	SDR	Energy	Food	Essentials
Treasury bill	1.0	5.3	27.4	12.2	5.8
UST 5-year	6.2	7.2	28.8	14.4	8.9
UST 10-year	11.0	11.4	31.3	17.4	13.2
UST 15-year	16.0	16.1	34.4	21.2	17.9
Gold	17.1	16.6	30.5	18.3	16.5
U.K. 10-year	12.8	10.6	30.0	17.3	13.9
Germany 10-year	12.0	8.3	30.1	16.8	13.2
Japan 10-year	13.3	10.9	30.4	17.5	14.2

The change in measuring unit materially alters the observed risk of conventional reserve assets. The three-month Treasury bill has annualized volatility of approximately 1 percent in dollars but 27 percent when measured relative to the world energy price index. The latter is close to the volatility of the corresponding unhedged dollar position measured in the same unit. Since short-term Treasury returns have little contemporaneous covariance with energy prices, most of the variation in the liability remains in the liability-unit return.

The contrast with the SDR is informative. The Treasury bill's volatility rises from 1 percent in dollars to approximately 5 percent in SDR units, but the relative ordering of the reserve assets remains broadly similar. Moving between monetary units therefore changes measured risk substantially less than moving from a monetary unit to a volatile goods liability. The broader essentials-weighted import basket lies between these cases: its diversification across goods produces considerably lower liability volatility, and the corresponding increase in measured asset risk is more moderate.

The change in unit also affects diversification across assets. Figure 2 compares pairwise correlations among conventional reserve assets in dollars and in units of world energy.

In dollars, the asset menu contains substantial cross-sectional variation. Short- and long-duration government securities have different risk exposures, gold has relatively low correlations with government bonds, and foreign bond markets provide additional variation. In energy units, these differences are substantially reduced. Every transformed return contains

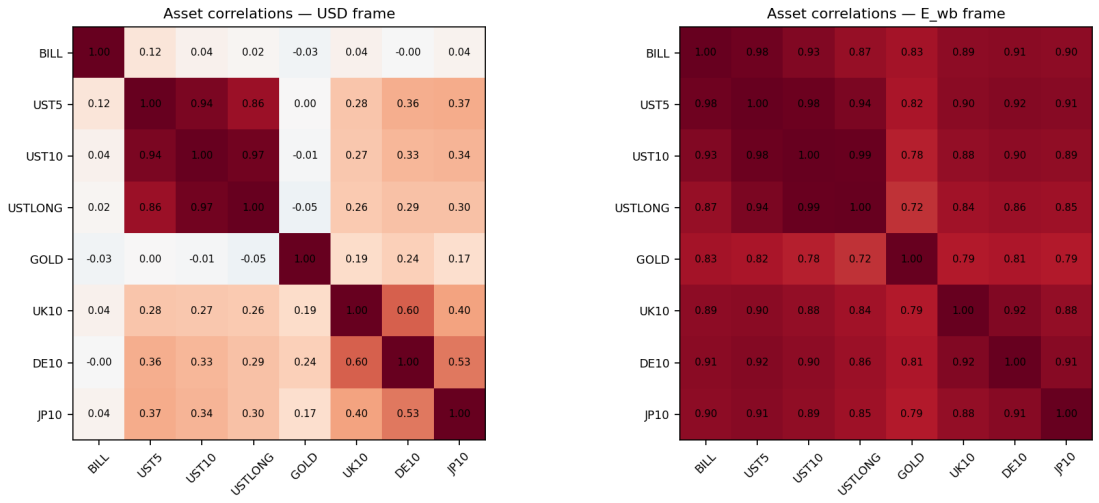


Figure 2: Pairwise correlations among conventional reserve assets in dollars and in units of world energy. In the energy frame, off-diagonal correlations range from 0.72 to 0.99, with a median of 0.89.

the same energy-price component in its denominator, and that common component is large relative to the variation in many of the underlying financial returns. As a result, pairwise correlations rise to between 0.72 and 0.99.

These results establish a measurement effect without yet establishing a portfolio effect. Conventional reserve assets that appear relatively safe and diversified in monetary units can exhibit substantially higher and more correlated risk when evaluated against an import-price liability. The next subsection asks whether reoptimizing the portfolio in that liability unit materially reduces the exposure.

5.2 The Hedging Deficit

We next solve the minimum-risk portfolio problem in Equation (2) separately for each measuring unit. This exercise distinguishes the effect of changing the unit of evaluation from the effect of changing portfolio weights within the conventional reserve opportunity set.

As a validation exercise, the same optimization procedure is first applied to the SDR. The resulting portfolio assigns substantial weight to the currencies represented in the SDR basket, even though the optimizer is provided only with the covariance matrix of transformed returns. The procedure therefore identifies liability-replicating combinations when the relevant exposures are present in the asset set.

The goods-based results are different. Figure 3 reports the minimum-risk weights across measuring units.

The dollar minimum-risk portfolio is approximately 99 percent Treasury bills. Reoptimization in goods units changes the allocation, but less than might be expected from the risk reclassification in Section 5.1. The portfolio evaluated against the world energy index remains approximately 82 percent in Treasury bills, with about 15 percent in gold and only small allocations to the remaining assets. The corresponding bill share is approximately

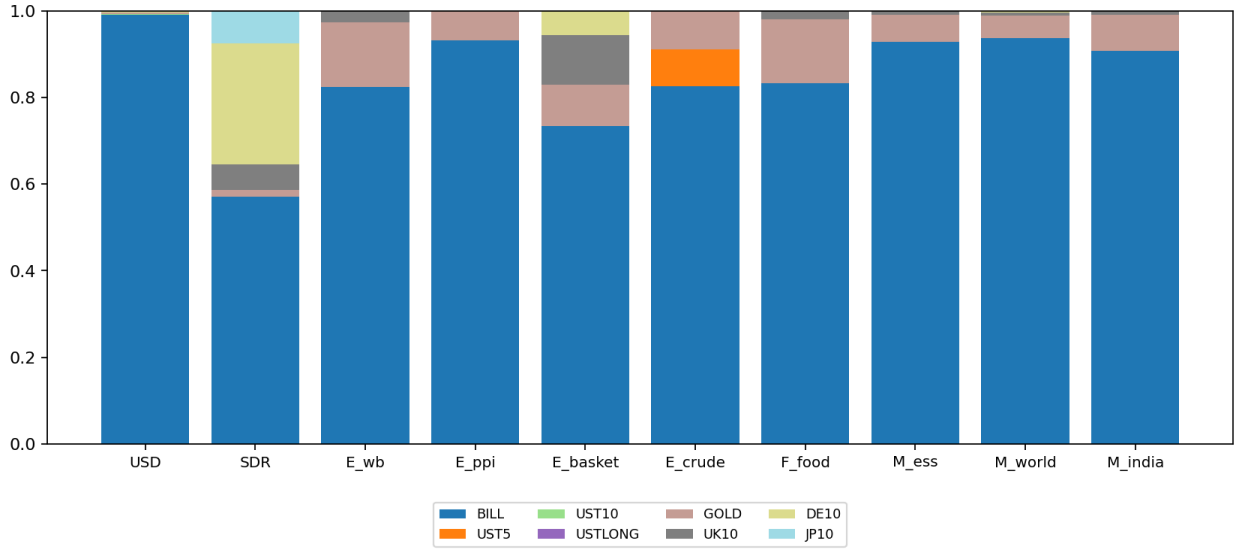


Figure 3: Minimum-risk portfolio weights by measuring unit. Portfolios are long-only, fully invested, and subject to the 30 percent Treasury-bill liquidity floor described in Section 4. Bars report point estimates; intervals are 90 percent moving-block bootstrap confidence intervals.

73 percent for the traded energy basket and 93 percent for the essentials-weighted import basket.

The high bill weights are not imposed by the liquidity constraint. The 30 percent floor is slack in the goods-based optimizations. Rather, the optimizer continues to favor short-duration assets because the other conventional reserve assets add financial-market volatility without providing sufficient positive covariance with the goods liability. The long-only restriction is more consequential: longer-duration bonds are negatively correlated with energy prices over much of the sample, so using that relationship to synthesize positive energy exposure would require short positions that are not part of the baseline reserve mandate. Additional diagnostics on these covariance relationships are reported in Appendix C.

Table 3 reports the resulting hedging deficit and hedgeable share for the two principal liability measures. At horizons beyond one month, volatilities are expressed throughout as per-month-equivalent annualized values—the volatility of compounded h-month returns rescaled to its one-month equivalent—so that magnitudes are comparable across horizons.

At the monthly horizon, the minimum-risk conventional portfolio removes approximately one-half of one percent of the variance associated with the world-energy liability. The hedging deficit, 26.1 percent annualized, is therefore close to the volatility of the corresponding unhedged dollar benchmark. The broader essentials basket is substantially less volatile in level terms, but the conventional portfolio also removes almost none of its variance at the monthly horizon.

The distinction between the **measurement effect** and the **allocation effect** is therefore important. Section 5.1 shows that changing the measuring unit materially alters asset volatilities and correlations. Table 3 shows that changing portfolio weights within the conventional reserve set does little to offset that change at short horizons. The principal limitation is

Table 3: Hedging deficit and hedgeable share. HD is the minimum attainable volatility in liability units, expressed as per-month-equivalent annualized percent. HS is the variance reduction relative to the zero-return dollar benchmark defined in Equation (4).

	World energy		Essentials basket	
Horizon	HD	HS	HD	HS
1 month	26.1	0.005	4.8	0.001
12 months	34.0	-0.055	7.9	0.020
24 months	32.1	-0.078	8.1	0.117
60 months	22.7	0.139	6.5	0.274

not that the initial portfolio was optimized under a monetary benchmark. It is that the conventional opportunity set contains little exposure that moves positively with the goods liability.

The horizon results lead to the same conclusion over the periods most relevant for reserve adequacy. Against world energy, the hedgeable share is slightly negative at twelve and twenty-four months. A negative value is possible because the zero-return dollar position used as the normalization is not itself a feasible yield-bearing reserve portfolio. At these horizons, the additional variation in feasible reserve returns slightly exceeds the amount of liability variation that their covariance offsets.

The result improves at longer horizons. At sixty months, the hedgeable share rises to approximately 0.14 for world energy and 0.27 for the essentials basket. The improvement in the energy specification is associated primarily with the stronger longer-horizon covariance between gold and energy prices. The sixty-month estimates are based on relatively few independent historical intervals, however, and are correspondingly less precise. The full term structure, including the food and alternative-energy specifications, is reported in Appendix C.

Overall, reoptimization within the conventional reserve asset set produces only limited reductions in goods-price risk over the horizons at which reserve adequacy is commonly assessed. This residual exposure is the hedging deficit defined in Section 3. The next subsection translates its magnitude from volatility units into import-cover units.

5.3 Implications for Reserve Adequacy

The preceding results describe residual risk in return units. We next express the same exposure using the Import-Months-at-Risk measure defined in Equation (6). Portfolios are re-estimated on rolling 120-month formation windows and evaluated over the subsequent twelve months, so the results in this subsection use only information available at the time each portfolio is formed.

Table 4 compares monetary volatility with the lower tail of import cover for portfolios optimized under different measuring units. Each import-cover calculation is normalized to an initial reserve position equal to three months of imports.

Measured against world energy prices, a reserve position beginning with three months of import cover has a fifth-percentile outcome of approximately 1.7 months after one year. This implies a lower-tail decline in energy purchasing power of more than 40 percent over

Table 4: Reserve-adequacy outcomes by portfolio mandate. Dollar volatility is realized out-of-sample, annualized; it therefore differs from the full-sample volatilities in Table 2. IMaR reports the fifth-percentile number of months of import cover remaining after twelve months from a standardized initial position of three months. Portfolios are estimated using rolling 120-month formation windows.

Mandate	Dollar volatility	IMaR: energy	IMaR: essentials
Dollar-optimal	0.75	1.73	2.58
SDR-optimal	4.58	1.72	2.64
Energy-optimal	4.73	1.75	2.62
Basket-optimal	4.85	1.67	2.47
Essentials-optimal	1.34	1.75	2.64
Bills only	0.74	1.73	2.58

the twelve-month horizon. Against the broader essentials-weighted basket, the corresponding fifth-percentile outcome is approximately 2.5 to 2.6 months.

More importantly for the portfolio exercise, the lower tail is similar across conventional reserve mandates. The dollar-optimal, SDR-optimal, and energy-optimal portfolios produce energy IMaR estimates between 1.72 and 1.75 months. The full range across the portfolios in Table 4 is 1.67 to 1.75 months. These differences are small relative to the differences in monetary volatility: estimated dollar volatility ranges from less than 1 percent for the bill-based and dollar-optimal portfolios to almost 5 percent for the basket-optimized portfolio.

The out-of-sample exercise therefore reproduces the allocation result from Section 5.2 in an economically more direct unit. Conventional reserve portfolios can differ substantially in their monetary risk characteristics while producing similar lower-tail outcomes for import purchasing power. Reoptimization within the conventional asset set changes the composition of the portfolio but has little effect on the downside energy-import capacity of the reserve stock.

The three-month initial position is a normalization rather than an assumption about the appropriate level of reserves. Because import cover evolves proportionally according to Equation (5), the reported quantities scale with alternative starting levels.

Appendix C reports the full import-cover distributions and the secondary certainty-equivalent comparison introduced in Section 3. The latter produces relatively small cross-numeraire welfare differences. At $\gamma = 5$, the estimated annual certainty-equivalent loss from holding the dollar-optimal rather than energy-optimal portfolio is approximately 17–29 basis points under the two expected-return specifications. This is consistent with the preceding results: the welfare comparison measures differences among portfolios inside the conventional opportunity set, while most of the goods-price exposure lies outside the span of that set.

The residual exposure is therefore economically material in the same units used to discuss reserve adequacy, even though the choice among conventional reserve portfolios has relatively little effect on its magnitude. We next examine whether this conclusion is different in adverse energy-supply states.

5.4 State Dependence

The unconditional results do not distinguish among the sources of energy-price movements. As discussed in Sections 2 and 4, this distinction may matter for the import-insurance interpretation of reserves. An increase in energy prices generated by strong global demand can coincide with different financial-market conditions from an increase generated by an adverse supply development. We therefore re-estimate the liability-unit covariance matrix and minimum-risk portfolio within the shock states defined in Section 4.

Table 5 reports the main conditional results. The baseline adverse supply state consists of observations in which the standardized price-raising oil-supply-news shock exceeds one-half of a standard deviation. A one-standard-deviation threshold provides a stricter specification. The table also reports selected demand-related states for comparison.

Table 5: Conditional hedging results in the world-energy frame. HD and the unhedged benchmark are per-month-equivalent annualized percent. Conditional covariance matrices are estimated over the post-1983 identification sample using the procedure described in Section 4.

State	n	Unhedged volatility	HD	HS	Gold weight
Supply, 0.5σ threshold	156	22.0	21.8	0.016	0.0
Supply, 1.0σ threshold	79	26.1	25.9	0.020	0.0
Complement	354	25.6	25.5	0.010	14.3
Demand–activity	69	24.7	23.6	0.086	0.0
Demand–consumption	77	36.8	36.0	0.041	47.0

The principal result is not that adverse supply states have unusually high unconditional volatility. Under the baseline classification, the volatility of the unhedged energy exposure is approximately 22 percent in supply months, compared with 26 percent in the complement. Some demand-related states exhibit substantially larger energy-price variation. The relevant question is instead whether conventional reserve assets provide greater covariance with energy prices when the supply state occurs.

They do not. The hedgeable share is 0.016 under the baseline supply-state definition and 0.020 under the stricter threshold. The corresponding hedging deficits remain very close to the unhedged benchmark. The highest hedgeable share among the states reported in Table 5 is 0.086 in the demand-activity state. Conditioning therefore changes the magnitude and composition of liability risk without materially changing the ability of the conventional reserve set to hedge it.

Gold provides a useful diagnostic because it is the conventional reserve asset most plausibly associated with adverse commodity-price states. Its unconditional correlation with world energy prices is positive but too small to offset its own volatility at short horizons. This conclusion is not reversed in the supply state. Under the baseline definition, gold's energy-price correlation is approximately 0.11, compared with a break-even value of approximately 0.25, and the minimum-risk supply-state portfolio assigns zero weight to gold. Under the more restrictive announcement-only identification, the corresponding correlation is approximately 0.15 against a break-even value of 0.20. The latter estimate is based on only 48 observations and is correspondingly imprecise, but the point estimate gives the same qualitative result.

Average returns, as distinct from covariances, do differ across states in one relevant way. Splitting oil-price increases by their identified source, gold is the only risky conventional reserve asset whose mean dollar return is higher in supply-driven months than in demand-driven months (approximately 1.25 versus 0.96 percent per month), whereas longer-duration bonds display the opposite pattern (approximately zero versus more than 1 percent per month). This asymmetry is consistent with gold providing some directional value in adverse supply states, but it does not alter the covariance-based conclusion: a favorable mean in a particular state is not a substitute for the co-movement required to reduce liability-unit variance.

The conditional estimates should be interpreted with two limitations in mind. First, the principal supply-state sample contains 156 observations, and confidence intervals on individual conditional correlations are wide. Second, the identifying series begins in 1983 without extrapolation. The oil shocks of 1973–74 and 1979–80 therefore motivate the historical discussion in Section 2 but do not enter the headline state-dependent estimates.

An additional ex-ante exercise examines the 2022 energy shock using portfolio weights estimated only with information available before the episode. A bills-only position beginning with three months of energy-import cover declines to approximately 2.1 months by mid-2022. The portfolio optimized ex ante in energy units reaches a slightly lower trough of approximately 2.0 months, reflecting its pre-event allocation to duration during a period in which interest rates and energy prices subsequently moved adversely together. A portfolio with a modest gold allocation produces a path close to the bills-only position. The exercise therefore provides an out-of-sample illustration of the same limitation identified in the covariance analysis: optimizing over the conventional reserve set does not materially alter the import-cover tail during the event. Additional event results, including the preliminary 2026 exercise, are reported in Appendix C.

Taken together, the results in this section establish four points. Changing the measuring unit materially alters the observed risk structure of conventional reserve assets. Reallocating across those assets removes little of the resulting goods-price exposure at short and medium horizons. The residual exposure is economically material when expressed in months of import cover. Finally, the limited hedgeability of the liability persists in adverse supply states and in ex-ante event exercises.

These findings motivate the next step of the analysis. Section 6 expands the opportunity set to instruments whose payoffs contain more direct exposure to energy prices and asks how much of the hedging deficit can be reduced once the conventional reserve-asset constraint is relaxed.

6 Extensions: Expanding the Sovereign Hedging Set

Section 5 shows that the conventional reserve asset set provides little protection against the goods-price exposure identified in the paper. This section asks how much of that result reflects the restriction of the opportunity set itself. We consider two extensions. First, we add financial instruments whose returns have more direct exposure to energy prices. Second, we broaden the unit of analysis beyond the central-bank reserve portfolio to include physical strategic stocks held elsewhere on the sovereign balance sheet.

The purpose of these exercises is comparative rather than prescriptive. Commodity-linked financial instruments and physical inventories differ from conventional reserves in liquidity, market depth, operational requirements, and institutional mandate. They are therefore not treated as direct substitutes for the entire reserve portfolio. Instead, the exercises measure how the attainable hedging set changes when instruments with progressively closer exposure to the underlying liability are introduced.

6.1 Commodity-Linked Financial Instruments

We begin by extending the baseline opportunity set to financial instruments with direct or indirect exposure to energy prices. The exercise is conducted on the common 1999–2024 sample available for the additional instruments. Commodity-futures positions are fully collateralized, so their returns consist of the Treasury-bill return plus the excess return from the rolled futures position. The same long-only, fully invested portfolio problem and 30 percent Treasury-bill floor used in the baseline analysis are then applied to the expanded asset set.

Figure 4 summarizes the results as the opportunity set is expanded sequentially.

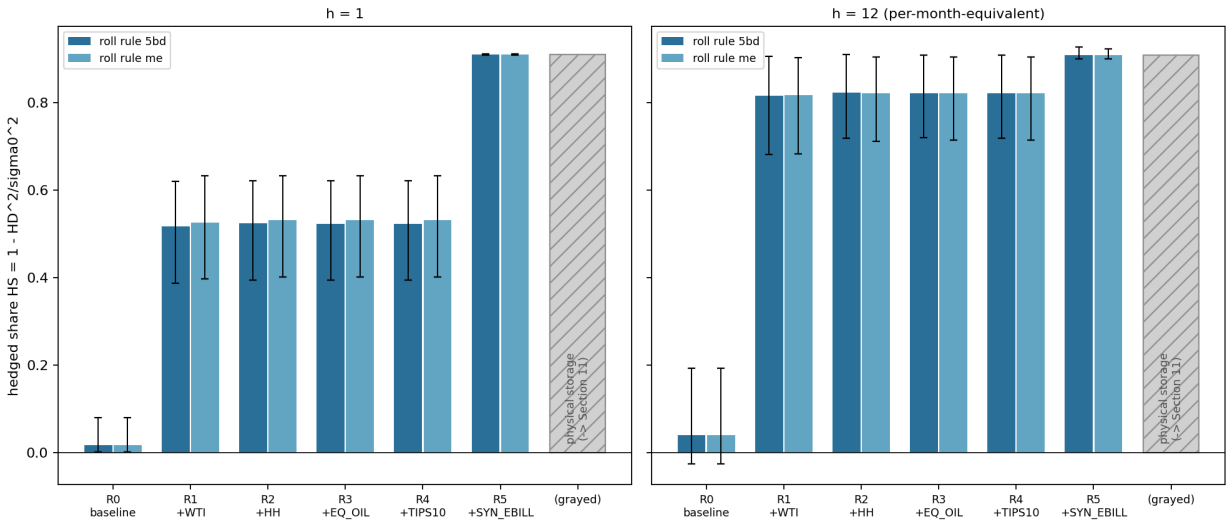


Figure 4: Hedgeable share as the financial opportunity set is expanded. Results are reported for the world-energy liability at one- and twelve-month horizons. Commodity-futures positions are fully collateralized. Bootstrap confidence intervals and alternative roll conventions are reported in Appendix D.

The addition of crude-oil futures produces the largest change. On the common sample, the conventional reserve menu hedges approximately 2 percent of the world-energy liability variance at the monthly horizon. Adding a rolled WTI futures position raises the hedgeable share to approximately 0.52. At the twelve-month horizon, the corresponding increase is from approximately 0.04 to 0.82. The optimizer assigns substantial weight to the futures exposure: approximately 53 percent of the portfolio at the monthly horizon and the maximum permitted 70 percent at twelve months.

The result is useful for interpreting the baseline findings. The optimization procedure is able to reduce the hedging deficit substantially once the asset set contains a payoff with direct positive exposure to the liability. The limited hedgeability documented in Section 5 therefore does not arise solely from the choice of portfolio criterion or from the transformation into liability units. It reflects the covariance structure of the conventional reserve asset set.

Additional financial instruments provide considerably smaller marginal improvements. Henry Hub natural-gas futures raise the monthly hedgeable share only slightly once WTI is already included. Energy-sector equities receive little or no weight in the energy specification after the futures positions are available, while U.S. inflation-linked government securities provide essentially no additional reduction in the variance of the world-energy liability. These results distinguish direct commodity-price exposure from more general inflation or energy-sector exposure. Detailed spanning tests and the corresponding results for the broader import basket are reported in Appendix D.

A more restrictive implementation illustrates the role of portfolio limits. If all non-baseline instruments are restricted to a 10 percent sleeve, the hedgeable share rises to approximately 0.19 at the monthly horizon and 0.23 at twelve months. The improvement is economically meaningful but substantially smaller than in the unconstrained extension. The difference reflects the size of the exposure selected by the liability-based optimizer rather than a change in the underlying covariance structure.

Taken together, these results establish that instruments with direct commodity-price exposure can materially expand the attainable hedging set. They do not imply that the statistical hedge is frictionless or that the resulting positions satisfy all of the institutional requirements placed on reserve assets. We turn to those limitations next.

6.2 Limits to Financial Replication

The futures-based results provide an implementable financial hedge, but not an exact claim on the liability used in the paper. The distinction can be measured by comparing the traded instruments with a frictionless benchmark.

We construct a hypothetical energy-indexed bill whose return equals the Treasury-bill return plus the change in the world energy price index. The instrument is not intended as a proposed security design. It provides a spanning benchmark: if a liquid financial claim moved one-for-one with the liability, how much of the hedging deficit could be eliminated under the same portfolio restrictions?

Under this benchmark, the hedgeable share rises to approximately 0.91 at both the monthly and twelve-month horizons. The remaining exposure is largely mechanical. The 30 percent Treasury-bill floor requires part of the portfolio to remain in a dollar asset that does not track the energy liability. Thus even an exact liability-linked claim does not eliminate all purchasing-power variance under the maintained reserve constraints.

The comparison with traded futures isolates a second source of residual risk. At the monthly horizon, the WTI-based portfolio reaches a hedgeable share of approximately 0.52 compared with 0.91 for the frictionless benchmark. At twelve months, the corresponding values are approximately 0.82 and 0.91. The difference is much larger at the monthly horizon than at the annual horizon.

Two frictions account for most of this gap. The first is basis risk. WTI is a crude-oil contract, while the headline liability is a broader world energy price index containing crude oil, natural gas, and coal. Over the common sample, the WTI exposure is highly correlated with the world-energy liability but still exhibits substantial tracking error. Changes in regional crude-price differentials and variation in natural-gas and coal prices therefore remain unhedged.

The second friction is the cost of maintaining futures exposure through time. The empirical returns used above already incorporate the realized gain or loss from rolling the contracts. These costs vary substantially across market regimes. WTI roll returns were favorable during periods of persistent backwardation but materially negative during the contango-dominated 2000s and 2010s. The full decade-by-decade calculations and alternative roll conventions are reported in Appendix D.

The effectiveness of the traded hedge also varies across states. Re-estimating the expanded portfolio within the adverse supply state raises the hedgeable share substantially relative to the conventional menu, but less than in the unconditional sample. On the available common state-dependent sample, the WTI-based extension produces a hedgeable share of approximately 0.28, compared with roughly 0.03 for the conventional asset set. The reduction in effectiveness is consistent with a wider basis between the traded crude contract and the broader energy liability during supply disruptions.

The financial extension therefore produces two related conclusions. First, direct commodity exposure substantially reduces the hedging deficit, confirming that the conventional asset menu is an important source of the baseline result. Second, the difference between a traded futures hedge and the frictionless liability-linked benchmark remains material because the financial contract does not exactly match the underlying import exposure and must be maintained through markets over time.

These limitations motivate a second extension. A sovereign can hold the underlying good itself rather than a financial claim whose payoff is linked imperfectly to its price.

6.3 Physical Strategic Reserves

Many governments maintain strategic inventories of energy and other essential goods outside the central-bank balance sheet. These stocks are institutionally distinct from foreign-exchange reserves, but they are relevant to the same external-purchasing-power problem considered in this paper. A stored unit of the imported good provides direct access to that good during a disruption and does not require the sovereign to replicate the exposure through a financial contract.

We incorporate stored crude into the framework as a separate asset whose return reflects the change in the crude price net of storage costs. Because the headline energy liability is broader than crude oil alone, the physical position still contains basis risk relative to the world energy index. It also entails carrying costs and requires storage capacity established before the shock occurs.

At the twelve-month horizon, replacing the futures exposure with stored crude raises the estimated hedgeable share from approximately 0.82 to 0.88. The corresponding frictionless liability-linked benchmark is approximately 0.91. Most of the remaining difference reflects the mismatch between crude oil and the broader energy basket rather than the assumed

storage charge: changing the storage-cost assumption affects expected returns but has comparatively little effect on covariance with the liability.

The economic comparison with financial reserves is more transparent when expressed in import-cover units. We therefore consider an illustrative sovereign beginning with three months of energy-import cover. A fraction s of that position is represented by physical crude and the remainder by the conventional financial reserve portfolio, with the physical share capped at one-half to reflect storage-capacity constraints.

Under the historical twelve-month evaluation windows, increasing the physical share improves the lower tail of energy-import cover throughout the permitted range. With the entire position held in financial reserves, the fifth-percentile outcome is approximately 1.73 months. With one-half of the initial cover represented by physical crude and storage costs calibrated to sovereign cavern operations, the corresponding outcome rises to approximately 2.30 months. The improvement remains present when the evaluation is restricted to adverse supply states, where the corresponding lower-tail measure rises from approximately 1.62 to 2.28 months.

This exercise is intentionally illustrative. It does not solve jointly for an optimal reserve stock, storage capacity, release rule, and financial portfolio. Its purpose is to establish that the attainable downside distribution changes once the sovereign opportunity set includes a direct physical claim on the imported good. Detailed assumptions regarding storage costs and alternative physical shares are reported in Appendix D.

Physical inventories also differ from financial hedges in ways that the portfolio statistic captures only partially. In the framework above, stored crude is valued primarily as protection against changes in the price of the good. In practice, a strategic stock can also provide quantity insurance when market access or physical availability is impaired. Conversely, physical inventories are less liquid than financial reserves for intervention, debt service, and other monetary uses. Their effectiveness depends on storage infrastructure, product grade, transportation capacity, and the rules governing release. A crude-oil stock, for example, does not eliminate exposure to refined-product spreads or to the natural-gas and coal components of the broader energy basket.

The relevant comparison is therefore not between two interchangeable reserve assets. Conventional financial reserves, commodity-linked claims, and physical inventories hedge different margins of external stress. Financial reserves provide liquid claims that can be mobilized for intervention and external payments. Commodity derivatives can reduce the price exposure associated with selected imports but introduce basis, rollover, and operational requirements. Physical stocks provide more direct protection against the stored commodity but require advance investment in inventories and infrastructure.

The extensions nevertheless sharpen the interpretation of Section 5. The large hedging deficit of the conventional reserve portfolio is substantially reduced only when the opportunity set includes instruments with direct exposure to the underlying liability. Traded financial contracts provide much of that exposure but not all of it; physical inventories provide an additional form of protection outside the conventional reserve portfolio. The next section examines the robustness of the main results and several additional considerations that are not required for this central comparison.

7 Robustness and Additional Considerations

The preceding sections rely on particular choices regarding the measurement of the import liability, construction of reserve-asset returns, portfolio restrictions, estimation, and identification of energy-price states. This section assesses whether the main conclusions are sensitive to those choices. We then consider a separate dimension of reserve risk that is not captured by the baseline analysis: the state-contingent usability of reserve assets.

The robustness exercises are organized around the principal result of the paper rather than around individual specifications. The relevant question is whether reasonable alternatives reveal substantial hedging capacity within the conventional reserve asset set. Across the exercises below, the estimated level of liability risk varies, but the answer to that question is largely unchanged. Full results are reported in Appendix E.

7.1 Measurement, Asset Construction, and Portfolio Constraints

We first consider whether the low hedgeable share in Section 5 depends on the construction of either the liability or the reserve asset set.

The result is insensitive to alternative measures of the energy liability. Replacing the headline world energy index with alternative fixed-weight energy baskets or an oil-only specification leaves the monthly hedgeable share between approximately 0.005 and 0.006. The level of the hedging deficit varies with the volatility of the particular price index, but the conventional reserve portfolio continues to eliminate little of that variance. The broader-import results are similarly insensitive to replacing the interpolated manufactured-goods component of the essentials basket with the available monthly non-fuel import-price series.

A country-specific exercise provides a more demanding measurement check. Using Türkiye's published import unit-value index rather than a constructed common liability measure produces a hedgeable share of approximately -0.01 over the available national-data sample. The bills-only twelve-month Import-Months-at-Risk measure is approximately 2.3 months from a standardized three-month initial position. The national index is only moderately correlated with the constructed import baskets, so this exercise changes the measured liability materially while leaving the portfolio conclusion unchanged. The shorter sample available for the national series limits the precision and historical coverage of the exercise, and the result is therefore treated as an external-validity check rather than an alternative baseline.

Alternative construction of the financial returns also has little effect on the results. Replacing the zero-coupon Treasury returns used in the headline analysis with par-coupon returns changes the goods-frame hedging deficit by at most approximately 0.03 percentage points. The timing diagnostic described in Section 4 moves the energy hedging deficit in the opposite direction from the main concern: using the lower-frequency specification raises the estimated deficit relative to the monthly estimate. The monthly-average construction of several goods-price series therefore appears, if anything, to attenuate the measured exposure.

Currency treatment of foreign government bonds is a particularly relevant alternative. Hedging the foreign bond positions back into dollars reduces their dollar volatility substantially, as expected, but has little effect on the world-energy hedging deficit. On the common comparison sample, the deficit is approximately 29.7 percent with unhedged foreign bonds and 29.9 percent with dollar-hedged bonds. Currency hedging changes the monetary expo-

sure of the instruments but does not create additional covariance with energy prices. The contrast is useful because the same operation materially worsens replication of the SDR, for which foreign-currency exposure is itself part of the liability. The two cases reinforce the distinction between currency matching and goods-price hedging developed in Section 2.

Finally, the baseline result does not appear to be generated by the portfolio constraints. The 30 percent Treasury-bill floor is slack in the goods-based minimum-risk portfolios, and imposing additional issuer caps has little effect on the headline hedging statistics. Relaxing or tightening these internal portfolio restrictions therefore does not reveal a conventional asset allocation with materially greater exposure to the goods liability.

Table 6 summarizes the main measurement and construction exercises.

Table 6: Selected robustness exercises. Reported alternatives are chosen to address economically distinct features of the baseline design. Full specifications and confidence intervals are reported in Appendix E.

Dimension	Alternative	Main implication
Energy liability	Alternative weights and oil-only specifications	Monthly HS remains approximately 0.005–0.006
Import liability	Türkiye national import unit-value index	HS approximately −0.01
Treasury returns	Par-coupon rather than zero-coupon construction	HD changes by at most approximately 0.03 pp
Price timing	Three-month diagnostic	Estimated energy HD increases
Foreign bonds	Currency-hedged rather than unhedged	Energy-frame HD essentially unchanged
Portfolio constraints	Non-U.S. issuer caps and alternative internal restrictions	No material change in goods-frame results

These exercises change the measurement of risk more readily than they change the ability of the conventional asset set to hedge it. The distinction is consistent with the main result of Section 5: the magnitude of the liability can depend on how that liability is defined, while the variance reduction available from conventional reserve assets remains limited.

7.2 Estimation, Samples, and State Identification

We next examine whether the results depend on estimation choices or on particular historical episodes.

The covariance estimator has negligible effects on the headline result. The annualized world-energy hedging deficit is approximately 26.15 percent using Ledoit–Wolf shrinkage and 26.11 percent using the sample covariance matrix. Alternative moving-block bootstrap lengths change the reported confidence bands but not the point estimates or qualitative inference. The central result is therefore not driven by covariance regularization.

The low hedgeable share also appears across distinct historical periods. Dividing the sample into four broad subperiods produces world-energy hedgeable shares of approximately 0.02, 0.09, 0.07, and 0.00. The estimated optimal composition changes substantially across those periods, particularly the allocation to gold, but no subperiod reveals a conventional reserve portfolio that hedges a large share of the liability. This is relevant because it distinguishes instability in estimated portfolio weights from instability in the underlying conclusion about attainable hedging capacity.

The identification of adverse supply states requires greater attention because the conditional analysis in Section 5.4 depends on a structural classification. Table 7 reports six alternative definitions. These include the baseline signed supply-news rule at two thresholds, an unsigned rule, direct classification using the structural oil-shock taxonomy, a Ready-style implementation, and the announcement-only version of the supply-news measure.

Table 7: Alternative definitions of the adverse supply state. HD is the world-energy hedging deficit and HS is the corresponding hedgeable share. The headline conditional sample excludes pre-1983 observations that require extrapolation of the identifying series.

Supply-state definition	n	HD	HS
Signed news, 0.5σ threshold	156	21.8	0.016
Signed news, 1.0σ threshold	79	25.9	0.020
Unsigned, $\lvert z\rvert > 0.5\sigma$	307	33.0	0.019
Structural taxonomy	112	20.7	0.014
Ready-style implementation	123	16.7	0.022
Announcement-only	48	29.6	0.036

The definition of the state affects the level of realized energy-price volatility and therefore the level of the hedging deficit. Across the six classifications, HD ranges from approximately 17 to 33 percent annualized. The normalized result is considerably more stable: the hedgeable share remains between 0.014 and 0.036. Alternative identification therefore changes which observations enter the supply state without revealing materially greater hedging capacity in the conventional asset menu.

The published pre-pandemic vintage of the supply-news series gives a similar result. State membership agrees with the later vintage for approximately 89–92 percent of observations under the relevant classifications, and the resulting conditional hedging deficit differs by only a small amount. We continue to exclude the extrapolated pre-1983 observations from the headline conditional estimates because their identification is qualitatively weaker.

The reserve-adequacy result is also not driven primarily by historical average returns. Repeating the Import-Months-at-Risk exercise after demeaning the underlying joint return paths lowers the fifth-percentile energy-cover outcome from approximately 1.73 to 1.63 months. Removing sample drift therefore changes the level modestly but leaves the lower-tail deterioration economically large.

Taken together, the estimation and sample exercises reinforce a distinction that recurs throughout the paper. Optimal weights can be unstable across samples and classifications, and the volatility of the liability can vary substantially across states. Neither form of variation produces a conventional portfolio with a large hedgeable share at the short and medium

horizons central to the analysis.

7.3 Reserve Usability and Sanctions Risk

The baseline framework assumes that a reserve asset can be mobilized at its market value when required. This assumption can fail independently of the relative-price exposure studied in the paper. Financial sanctions, asset freezes, custody arrangements, settlement restrictions, and related measures can make a reserve asset partially or fully unavailable to its holder even when its market price is unchanged.

This form of risk can be represented conceptually as a state-contingent usability haircut on reserve assets. If $a_{i,s} \in [0, 1]$ denotes the fraction of asset i's market value that is accessible in state s, then the effective value of the reserve position in that state is $a_{i,s}V_i$ rather than V_i. The sanctions literature studies both the determinants of such exposure and the resulting changes in reserve composition, including increased gold holdings by central banks facing greater risk of financial sanctions by reserve-issuing jurisdictions (Arslanalp et al., 2023; Ferranti, 2025).

Usability risk is distinct from the liability-price risk estimated in Sections 5 and 6. The latter asks whether the payoff on an accessible reserve asset co-moves with the price of the import liability. Sanctions risk asks whether the asset can be accessed in the first place. An asset can therefore perform differently along the two dimensions. For example, greater holdings of gold can be consistent with protection against reserve-freeze risk without implying that gold provides a short-horizon hedge against energy prices.

Custody introduces an additional trade-off. Assets held outside the jurisdiction of a potential sanctioning authority may be less exposed to legal restrictions, but decentralized or domestic custody can reduce immediate market liquidity or increase the operational cost of mobilization. Gold illustrates the point particularly clearly: the location and legal status of the asset affect its accessibility, while its market covariance with an import liability is a separate empirical object (Arslanalp et al., 2023; McDowell, 2021).

The interaction between sanctions and commodity-price shocks can be economically important when the states occur simultaneously, as in parts of the 2022 episode. The present analysis does not estimate a joint distribution of supply shocks and reserve-access restrictions, and no structural relation between the two is assumed. A combined framework would require modeling both the payoff of reserve assets relative to the import liability and their state-contingent usability.

This distinction places a useful boundary around the results of the paper. The finding that conventional reserve assets provide little hedge against energy-price movements does not imply that their other reserve-management characteristics are unimportant. Conversely, evidence that a particular asset provides protection against sanctions or custody risk does not establish that it hedges import-price risk. The two exposures imply separate portfolio considerations.

7.4 Scope and Limitations

Several limitations apply to the empirical exercises.

First, the headline goods liabilities are representative rather than universal. The relevant stress-period import basket is country-specific, and internationally comparable monthly import-price data are limited. The country-specific Türkiye exercise provides one check using a national index, but its available sample begins substantially later than the common historical panel.

Second, the state-dependent analysis is constrained by identification. The preferred supply-news measure is available without extrapolation only from 1983 onward. The major oil disruptions of the 1970s therefore inform the motivation for the exercise but do not enter the headline conditional estimates. Conditional covariance estimates are also less precise than the unconditional results because the relevant state samples are smaller.

Third, estimates at the longest horizons are based on relatively few independent historical intervals. The positive five-year hedgeable share associated with longer-horizon gold covariance should therefore be interpreted with greater uncertainty than the monthly and annual results.

Fourth, some data checks cannot be completed using freely redistributable historical series. A separate bank-deposit return leg would require reconstruction of a discontinued benchmark, while validation of the constructed government-bond returns against long commercial total-return histories requires licensed data. Neither limitation affects the asset classes or liability measures driving the headline results, but both remain useful external checks.

Fifth, the extensions in Section 6 are partial-equilibrium exercises. They treat the reserve manager or sovereign as taking market prices, storage capacity, and financial-market depth as given. Large coordinated shifts by official reserve holders toward commodity-linked claims could affect the prices and availability of the same instruments being evaluated. The physical-reserve calculation likewise abstracts from endogenous investment in storage capacity and from optimal release and replenishment rules.

Finally, observations from 2026 are used only in explicitly identified out-of-sample or descriptive exercises and remain preliminary. They are not included in the baseline estimation sample.

The robustness exercises therefore support a relatively narrow conclusion. The measured level of goods-price risk varies across liability definitions, samples, and states, but conventional reserve assets eliminate only a small share of that risk at short and medium horizons under a wide range of specifications. This conclusion concerns the purchasing-power dimension of reserve risk. Liquidity, intervention capacity, debt-service needs, credit risk, and state-contingent asset usability remain distinct objectives and constraints of reserve management.

8 Implications for Reserve Management

The results have several implications for the interpretation of reserve management. They do not imply that a single measuring unit is appropriate for the reserve portfolio as a whole. International reserves serve multiple functions, including foreign-exchange intervention, external debt service, liquidity provision, and precautionary protection against disruptions to external financing. Monetary numeraires remain natural for many of these purposes. The

narrower implication of the analysis is that the relevant measure of risk depends in part on the liability associated with a particular reserve function. For the component of reserves intended to preserve the capacity to finance imports, monetary risk measures and liability-unit risk measures describe different dimensions of the same portfolio. Appendix F summarizes this mapping between reserve functions and their associated risk objects.

This distinction is visible in the comparison between monetary portfolio risk and reserve adequacy. Across the conventional portfolios considered in Section 5, dollar volatility varies substantially while the lower tail of energy-import cover changes very little. Portfolios that differ markedly in duration, currency exposure, and monetary volatility can therefore generate similar outcomes when evaluated in terms of the quantity of imports they can finance. Conversely, a portfolio that is stable in its reporting currency can remain exposed to large movements in import purchasing power. Portfolio risk and reserve adequacy are consequently related but distinct objects.

The hedging deficit provides one way to separate these objects. A high liability-unit volatility need not indicate that the reserve portfolio has been inefficiently constructed. It can instead reflect variation in the external liability that cannot be materially reduced using the assets available under the reserve mandate. This interpretation is particularly relevant to the results in Section 5, where reoptimization across conventional reserve assets changes the portfolio composition but eliminates only a small share of energy-price variance. In this setting, the residual exposure is primarily a property of the interaction between the liability and the feasible asset set rather than of a particular choice of weights within that set.

The Import-Months-at-Risk exercise expresses the same distinction in reserve-adequacy units. Traditional import-cover measures describe the initial size of the reserve stock relative to imports. The analysis here adds a stochastic component by allowing both asset values and import prices to evolve over the measurement horizon. An economy beginning with a given level of import cover can therefore face different distributions of future cover depending on the volatility of its import prices and their covariance with its reserve assets. This does not alter the interpretation of conventional reserve-adequacy measures; it identifies an additional dimension of the purchasing power represented by a given initial reserve stock.

The expanded opportunity sets in Section 6 further clarify the role of portfolio feasibility. When commodity-linked financial instruments are added, the attainable reduction in energy-price risk increases substantially. The comparison with the hypothetical liability-indexed bill shows that part of the remaining exposure can be attributed to basis risk, rollover costs, and reserve-portfolio constraints rather than to the portfolio method itself. These exercises therefore distinguish the absence of suitable covariance in the conventional reserve menu from the implementation frictions associated with obtaining that covariance through traded financial instruments. They do not establish an optimal allocation to those instruments, which would require additional assumptions about liquidity needs, operational capacity, expected returns, and institutional mandates.

Physical strategic stocks represent a related but institutionally distinct margin. Foreign-exchange reserves and strategic commodity inventories are not interchangeable assets. The former provide liquid financial claims that can be mobilized for intervention, debt service, and a broad range of external payments; the latter provide direct access to particular goods but depend on storage, transportation, product specification, and release capacity. Nevertheless, the two can affect the same economic outcome considered in the adequacy exercise: the

quantity of essential imports that the sovereign can finance or supply during an external shock. The physical-stock results in Section 6 therefore identify a connection between balance sheets that are normally analyzed separately without implying that they belong within a common institutional mandate or accounting category.

The same separation of functions applies to the usability risks considered in Section 7. An asset's ability to preserve purchasing power over imports is distinct from its accessibility under sanctions, custody restrictions, or settlement disruptions. Gold, for example, can have properties that are relevant to reserve usability without providing substantial short-horizon covariance with energy prices. More generally, reserve assets can perform differently across liquidity risk, currency risk, credit risk, access risk, and import-price risk. The results in this paper concern the last of these and do not rank assets across the broader set of reserve-management objectives.

Taken together, the analysis places the choice of measuring unit, the feasible asset set, and the definition of reserve adequacy within a common framework. Monetary portfolio statistics characterize risks relevant to the financial functions of reserves. Liability-unit statistics characterize the purchasing power of those same assets relative to particular external obligations. Strategic physical stocks affect part of the same external exposure through a different institutional mechanism. These measures and instruments therefore describe complementary dimensions of sovereign external resilience rather than alternative definitions of a single optimal reserve portfolio.

The analysis does not determine the appropriate size of reserves, the appropriate share of any asset class, or the allocation of responsibilities across public institutions. Those questions depend on objectives, constraints, and welfare considerations outside the portfolio problem studied here. The contribution is instead to identify which component of import-price risk is captured by conventional reserve assets, which component remains unspanned, and how that distinction changes when the relevant opportunity set is broadened.

9 Conclusion

This paper examines conventional reserve portfolios in units of the external liabilities they may be required to finance. The analysis is motivated by a narrow feature of reserve adequacy: for the precautionary component of reserves associated with import financing, the relevant economic value of the reserve stock depends partly on the quantity of imports it can purchase. Evaluating reserve assets relative to that liability produces a substantially different risk structure from the one observed in conventional monetary units.

The principal result is that this change in measurement has much larger effects on measured risk than on the risk reduction attainable through conventional portfolio reallocation. In the world-energy specification, assets that are relatively safe and diversified in monetary units become substantially more volatile and highly correlated in liability units. Reoptimizing the portfolio within the conventional reserve asset set removes only a small fraction of this exposure at short and medium horizons. The same conclusion appears when the residual risk is expressed in months of import cover and when the analysis is conditioned on adverse energy-supply states. It is also robust to alternative liability measures, return constructions, portfolio restrictions, estimation methods, samples, and state definitions.

The expanded-asset exercises clarify the source of this result. When the opportunity set includes financial instruments with direct commodity-price exposure, the attainable reduction in liability risk increases substantially. The remaining gap relative to a frictionless liability-linked claim reflects basis risk, rollover costs, and maintained reserve constraints. Physical strategic stocks provide a further, institutionally distinct source of exposure to the underlying good. These results indicate that the hedging deficit identified in the conventional portfolio is primarily a limitation of the available instrument set rather than of the portfolio method itself.

The analysis does not characterize an optimal reserve allocation or institutional design. International reserves serve multiple monetary and precautionary functions, and the appropriate benchmark depends on the objective under consideration. The results instead identify a distinction between the financial risk of reserve assets and their purchasing power over particular external liabilities. More generally, the assessment of reserve adequacy depends not only on the size and monetary risk of the reserve stock, but also on the liabilities it is expected to finance and the instruments available to manage those exposures.

A Historical Episodes

This appendix briefly summarizes the principal historical episodes referred to in Section 2. The purpose is descriptive: each episode illustrates, under different financing conditions, the distinction between the monetary value of external assets and their purchasing power over essential imports. Figure 1 in the main text spans the same period.

1973–74. Following the October 1973 embargo and associated production cuts, the posted price of benchmark Saudi crude rose from $3.01 per barrel to $11.65 by January 1974. A dollar asset that retained its nominal value over those months lost roughly three-quarters of its purchasing power over oil on that price measure; posted and realized market prices differed during the transition (Adelman, 1990, Table VI).

1979–80. Iranian production fell by approximately 4.8 million barrels per day during the revolution, around 7 percent of world production, and spot crude prices more than doubled between April 1979 and April 1980, while posted official prices rose by less. Part of the increase reflected precautionary demand and inventory accumulation rather than lost supply alone, which is one motivation for the state classification used in the conditional analysis (Hamilton, 2013; Barsky and Kilian, 2002).

Sri Lanka, 2022. By end-2022, gross international reserves had fallen to approximately $1.9 billion, around one month of prospective imports, of which roughly $1.4 billion was a renminbi deposit that was not then usable. IMF staff described the remaining freely usable reserves as insufficient for essential imports including fuel. Panel B of Figure 1 shows the realized import-cover path through the default (International Monetary Fund, 2023b).

Euro area, 2022. The euro area's energy goods trade deficit increased from 1.9 percent of GDP in 2021 to 4.0 percent in 2022, and the current account moved from a 2.8 percent

surplus to a 0.8 percent deficit as then published; subsequent revisions place the 2022 current account at −0.2 percent of GDP. The adjustment operated through the terms of trade and external financing rather than through foreign-exchange rationing (European Central Bank, 2023; Eurostat, 2026).

Ghana, 2022–23. By February 2023, gross international reserves on the IMF programme definition, which excludes encumbered assets, had fallen to approximately $1.1 billion, equivalent to about half a month of imports; headline gross reserves including encumbered assets were near $5.9 billion on the same date. The authorities' Gold-for-Oil programme used domestically purchased gold, directly or through sale, to finance petroleum imports; programme documents later identified associated balance-sheet risks, and the programme was discontinued in 2025 (International Monetary Fund, 2025; Bank of Ghana, 2026).

Hormuz, 2026. The 2026 disruption of shipping through the Strait of Hormuz lies outside the estimation sample and enters only the out-of-sample exercises. Japan announced the release of approximately 80 million barrels from its petroleum reserves in March, within IEA-confirmed member commitments totaling 412 million barrels, including 172 million barrels authorized from the U.S. Strategic Petroleum Reserve. Japan also undertook coordinated foreign-exchange intervention with the United States in July, and U.S. officials discussed expanded use of the Federal Reserve's FIMA repo facility, which provides dollar liquidity against Treasury collateral. The episode is notable for combining physical releases, currency intervention, and collateralized dollar liquidity within a single external-stress response (Ministry of Economy, Trade and Industry, Japan, 2026; International Energy Agency, 2026; Ministry of Finance, Japan, 2026; U.S. Department of the Treasury, 2026). All 2026 figures are preliminary.

B Data Construction and Empirical Implementation

This appendix documents the financial return series, goods-price measures, sample construction, state classification, portfolio restrictions, and estimation procedures used in the empirical analysis. The objective is to make the construction of every headline series reproducible while keeping series-specific details out of the main text.

B.1 Financial asset returns

B.1.1 U.S. government securities

The U.S. asset set contains Treasury bills and intermediate- and long-duration Treasury securities. Bill returns are constructed from secondary-market discount rates converted to bond-equivalent yields. Longer-duration returns are constructed from the fitted Treasury yield curve of Gürkaynak et al. (2007).

Constant-maturity zero-coupon returns provide the baseline construction for the main exhibits. Par-coupon returns generated from the same yield curves are retained as an alternative specification because they more closely approximate securities held in actual reserve

portfolios. The longest maturity used in the historical sample is fifteen years, reflecting the maturity coverage of the fitted U.S. curve at the beginning of the sample.

For the opening portion of the sample in which the fitted curve is unavailable at all required maturities, returns are constructed from constant-maturity yields using the duration approximation described below. All such observations are identified in the series inventory.

B.1.2 Foreign government bonds

German, Japanese, and U.K. government-bond returns are constructed from official yield series using the same return approximation. The baseline specification leaves the currency exposure on these bonds unhedged. This treatment is consistent with an exercise in which currency composition is itself part of the reserve-allocation decision. A separate robustness exercise converts foreign-bond returns into currency-hedged dollar returns before applying the liability-unit transformation (Campbell et al., 2010).

The German and U.K. yield series begin in early 1971. Market-based Japanese government-bond yields become available in September 1974, creating the binding start date for the full reserve-asset panel: the joint sample begins in October 1974, the first month in which returns are available for every asset in the conventional menu.

The early Japanese series should also be interpreted with some caution because the JGB market remained relatively thin and partially administered during part of the 1970s and early 1980s. The subsample analysis in Appendix E, which reports the pre-1985 era separately, bounds the influence of this segment.

B.1.3 Bank deposits

No separate bank-deposit asset is included in the baseline menu. Deposits and short-term government bills provide closely related short-duration monetary exposure for the purpose of the goods-unit exercise, while a continuous historical deposit benchmark would require a splice after the discontinuation of the relevant eurodollar series. A separate deposit leg is left as a data extension rather than as part of the baseline construction.

B.1.4 Gold

Gold prices are taken from the London afternoon fixing as distributed through the World Bank commodity database (World Bank, 2026). The freely available historical series is reported as a monthly average of daily observations rather than an end-of-month price.

Because monthly averaging reduces measured volatility relative to point-in-time observations, the baseline series is accompanied by timing diagnostics using available market-price data over the modern sample. Exchange-traded gold futures provide an additional cross-check beginning in the mid-1970s (Erb and Harvey, 2013); the averaging adjustment follows Working (1960).

The baseline analysis uses the published monthly series without mechanically rescaling it. Any end-of-month-equivalent volatility calculations are reported as diagnostics rather than substituted for observed data.

B.1.5 SDR

The SDR is valued using the IMF's official dollars-per-SDR exchange rate. The monthly series extends through both the earlier administered regime and the subsequent currency-basket regime. The SDR is used as a diversified monetary control.

B.1.6 Commercial-index validation

Where commercial total-return indices are available, they may be used to validate the constructed government-bond return series over their overlapping sample. Because these data require subscriptions, they are not used to generate headline estimates or any result required for replication. The freely available constructed series remains the basis of the paper.

B.2 Goods-price and import-price measures

B.2.1 Energy

The baseline energy liability is the World Bank world energy price index, available monthly over the full empirical sample (World Bank, 2026).

The world index is preferred to domestic producer or consumer energy-price indices because the object of interest is the external price faced by an importer at or near the border. Domestic energy-price measures incorporate taxation, regulation, distribution margins, subsidies, and incomplete pass-through. U.S. producer fuel-price indices are therefore used only as alternative liability definitions.

A second energy measure is constructed from internationally traded prices for crude oil, natural gas, and coal. From 1990 onward, the components are converted to a common energy unit and combined using fixed world energy-consumption weights. The baseline weighting vintage uses current global primary-energy shares; a 1990-base weighting and alternative component weights provide robustness checks.

The underlying commodity series are taken from the World Bank commodity database. The European natural-gas series contains a documented change in pricing basis in 2015 from border-price measures toward hub pricing. We therefore report alternative basket specifications, including an oil-only variant unaffected by this change (Appendix E).

B.2.2 Food

The food liability uses the World Bank food price index, which is available over the full baseline sample. The FAO food price index provides an independent comparison beginning in 1990. Wheat, maize, and rice prices are retained for component-level diagnostics (World Bank, 2026; Food and Agriculture Organization, 2026).

B.2.3 Broader import and essentials baskets

The broader import-price measure combines energy, food, and a residual manufactured-goods component. The residual is proxied by the World Bank manufactures unit-value index. Because that index is available annually over much of the historical sample, it is

interpolated to monthly frequency for the baseline construction. The interpolation method is fixed before portfolio estimation and documented in the replication code.

A U.S. BLS non-fuel import-price index provides an alternative monthly residual beginning in the mid-1980s. Results using that series are reported over its own common sample rather than spliced into the earlier history.

The baseline basket uses weights intended to approximate essential-import exposure during external stress. A conventional average import basket is reported as an alternative. Import weights are based on World Bank WITS data and are held fixed at a pre-pandemic 2019 vintage so that portfolio results are not mechanically affected by contemporaneous crisis-driven changes in trade composition.

B.2.4 Country-specific measures

The headline analysis uses a representative importing-economy construction so that the same liability definitions can be evaluated over a long historical sample.

India provides a country-weighted application using the representative price indices with India-specific import shares. A continuous public monthly national import-price index is not used for the full sample.

Türkiye provides the principal validation using a nationally published import unit-value index. The currently accessible monthly national series begins in 2013. Türkiye is also informative because the import-price exercise in Beck and Rahbari produces one of the largest changes in their estimated optimal reserve composition (Beck and Rahbari, 2008, Table 6).

Sri Lanka is treated separately in the event analysis. Monthly reserve data and lower-frequency import data are used to construct realized import-cover paths through the 2022 crisis. The conversion of quarterly import flows into monthly rates and the definition of usable versus gross reserves are documented with the event-study files (Central Bank of Sri Lanka, 2026; International Monetary Fund, 2026b).

B.3 Sample construction

The master monthly sample runs from January 1971 through December 2025. The starting date provides a long post-Bretton-Woods sample in which major reserve assets and internationally traded goods prices can be evaluated under changing relative prices.

Individual asset and liability series have different start dates. Descriptive statistics may therefore use the longest valid common sample for each pair of series. Portfolio results use a common sample containing every asset required by the relevant specification.

The full conventional reserve menu is available from October 1974, the first month in which returns exist for every menu asset; that date is used consistently in the text, tables, figures, and replication files.

The principal sample conventions are:

- descriptive pairwise statistics: longest available overlapping sample;
- baseline conventional reserve portfolio: October 1974 through December 2025;

- conditional shock analysis: July 1983 through December 2025;
- expanded financial-menu analysis: common period available for all additional instruments, currently 1999–2024 for the full overlay exercise;
- 2026 observations: reserved for explicitly labelled out-of-sample applications.

Every table and figure records its effective estimation window.

B.4 Identification of energy-price states

B.4.1 Baseline oil-supply-news measure

The baseline supply-state classification uses the published oil-supply-news series of Känzig (2021). The analysis uses a fixed data vintage through December 2025.

The identification is based on revisions to expected future oil supply inferred from futures-price movements around OPEC-related news. Because the estimated structural series can change when the underlying sample is extended, the data vintage is fixed for the empirical analysis. Earlier published vintages are retained for sensitivity analysis.

The headline conditional estimates begin in July 1983. Pre-1983 observations rely on model-based extrapolation and are not used in the baseline conditional results.

The published shock is standardized before classification. The baseline adverse supply state contains observations in which the price-raising supply shock exceeds 0.5 standard deviations. A threshold of one standard deviation provides a stricter definition.

B.4.2 Announcement-only identification

The published announcement-surprise series is sparse because months without qualifying OPEC announcements contain no announcement shock. An announcement-only specification is therefore retained as a stricter identification exercise. It reduces the possibility that the same broader monthly oil-price movements used to define states also mechanically influence the covariance estimates.

The announcement-only specification is not used as the baseline because economically important supply news can occur outside the formal announcement calendar.

B.4.3 Demand-side classification

For observations outside the baseline adverse supply state, the structural oil-market decomposition of Baumeister and Hamilton is used to distinguish global-activity and oil-specific demand shocks where required (Baumeister and Hamilton, 2019).

The classification is used to determine whether reserve-asset covariance differs systematically across economically different sources of oil-price movements. It is not required for the unconditional portfolio analysis.

B.4.4 Alternative classifications

The robustness analysis includes:

1. a one-standard-deviation supply threshold;
2. an unsigned shock rule;
3. the announcement-only Känzig measure;
4. a partition based directly on the alternative structural shock taxonomy;
5. an implementation following the identification logic of Ready (2018);
6. earlier vintages of the supply-news series.

The agreement between alternative classifications and the corresponding conditional portfolio estimates are reported in the robustness appendix.

The 1973–74 and 1979–80 episodes predate the clean instrument period and are therefore not assigned headline conditional estimates. They are used only in historical or event-based interpretation.

B.5 Portfolio constraints

The baseline feasible set imposes three restrictions:

1. portfolio weights are non-negative;
2. portfolio weights sum to one;
3. at least 30 percent of the portfolio is held in U.S. Treasury bills.

The bill floor approximates the liquidity component of reserve portfolios documented in the institutional evidence discussed in Section 2. It is not intended as a universal regulatory constraint.

Issuer caps are excluded from the baseline. A robustness specification limits exposure to each non-U.S. sovereign issuer to 25 percent of the portfolio.

The expanded-menu exercises modify the asset set while otherwise retaining the portfolio framework. The specific treatment of futures collateral, roll costs, inflation-linked securities, commodity-related assets, and physical reserves is documented with the corresponding extension.

B.6 Covariance estimation and uncertainty

Headline covariance matrices are estimated using Ledoit-Wolf shrinkage (Ledoit and Wolf, 2004). Conventional sample covariance estimates are retained as a comparison.

Portfolio uncertainty is estimated using a moving-block bootstrap of the joint asset-and-liability panel. The baseline procedure uses:

- 24-month blocks;
- 1,000 bootstrap replications;
- 90 percent confidence intervals.

Block lengths of 12 and 36 months provide sensitivity checks.

For horizon $H > 1$, returns are compounded over overlapping H-month windows before the liability-unit transformation is applied. Bootstrap blocks are lengthened for the long-horizon exercises so that dependence within the constructed returns is not treated as independent sampling variation.

The principal horizons are one, twelve, twenty-four, and sixty months. Long-horizon portfolio weights are interpreted cautiously because the number of effectively independent observations declines substantially with the horizon.

B.7 Timing conventions

Financial asset prices are generally observed at or near the end of the month, while several commodity-price series are published as monthly averages. This creates a potential timing mismatch in the one-month covariance estimates.

The baseline uses the published frequency of each series rather than creating synthetic end-of-month commodity prices where those observations are unavailable. A three-month horizon provides a diagnostic because short differences in monthly timing are less important after aggregation.

Series based on monthly averages are identified in the data inventory. Additional timing comparisons using higher-frequency observations are reported where freely available.

B.8 Expected returns and the welfare comparison

Expected returns are required only for the cross-numeraire welfare calculation. They do not enter the minimum-variance portfolios, hedging deficit, or hedgeable share.

Two expected-return specifications are reported:

1. transformed sample means shrunk halfway toward the cross-asset grand mean;
2. a zero-excess-return specification in which expected excess returns relative to Treasury bills are set to zero.

The second specification removes cross-asset mean-return differences and therefore places the welfare comparison close to the covariance-based portfolio problem. Results are reported across the risk-aversion grid specified in the empirical code.

Alternative expected-return estimators are treated as sensitivity checks rather than as headline specifications.

B.9 Replication and data integrity

The empirical pipeline records the expected format and basic properties of each input series before estimation. Automated checks include date coverage, duplicate observations, missing-value patterns, and series-specific properties required by the state-identification data.

Each empirical output is linked to the script and panel version that generated it. The replication package records the data build date and code version for each table and figure.

Headline results are constructed from freely redistributable data. Licensed series, where available, are used only as validation checks and are not required to reproduce the principal findings.

C Additional Results for Section 5

This appendix reports additional results underlying Section 5. The material includes the full risk-reclassification grid, covariance diagnostics, alternative return constructions, the complete term structure of the hedging deficit, the secondary welfare comparison, and additional event exercises.

C.1 Full Risk Map and Return-Construction Checks

Table 8 reports the full volatility grid underlying Section 5.1. In addition to the baseline zero-coupon Treasury returns shown in the main text, the table reports par-coupon Treasury returns, the U.S. producer-fuel price variant, and the alternative traded-energy basket.

Table 8: Annualized asset volatility by measuring unit, monthly horizon. Columns match Figure 5; the producer-fuel column measures domestic U.S. fuel prices rather than the border price.

	USD	SDR	E (US prod.)	E (world)	E (basket)	Food	M (essentials)
Treasury bill	1.0	5.3	13.0	27.4	24.6	12.2	5.8
UST 5y (zero)	6.2	7.2	15.2	28.8	25.7	14.4	8.9
UST 10y (zero)	11.0	11.4	18.6	31.3	27.6	17.4	13.2
UST 15y (zero)	16.0	16.1	22.6	34.4	29.9	21.2	17.9
UST 5y (par)	5.4	6.6	14.8	28.6	25.6	14.0	8.3
UST 10y (par)	8.1	8.8	16.7	30.1	26.9	15.5	10.7
UST long (par)	10.0	10.5	18.2	31.3	28.0	16.8	12.4
Gold	17.1	16.6	20.1	30.5	26.5	18.3	16.5
U.K. 10y	12.8	10.6	18.3	30.0	26.0	17.3	13.9
Germany 10y	12.0	8.3	18.0	30.1	26.2	16.8	13.2
Japan 10y	13.3	10.9	19.4	30.4	27.3	17.5	14.2

The distinction between zero-coupon and par-coupon return construction is material for the dollar volatility of longer-duration Treasury securities but considerably less important in goods units. For example, dollar volatility at the longest U.S. maturity is 16.0 percent under the zero-coupon convention and 10.0 percent under the par-coupon convention, while the corresponding energy-unit volatilities are 34.4 and 31.3 percent. The common liability-price

component therefore accounts for a larger share of measured risk once returns are expressed in goods units.

The full heatmap corresponding to Table 8 is reported in Figure 5.

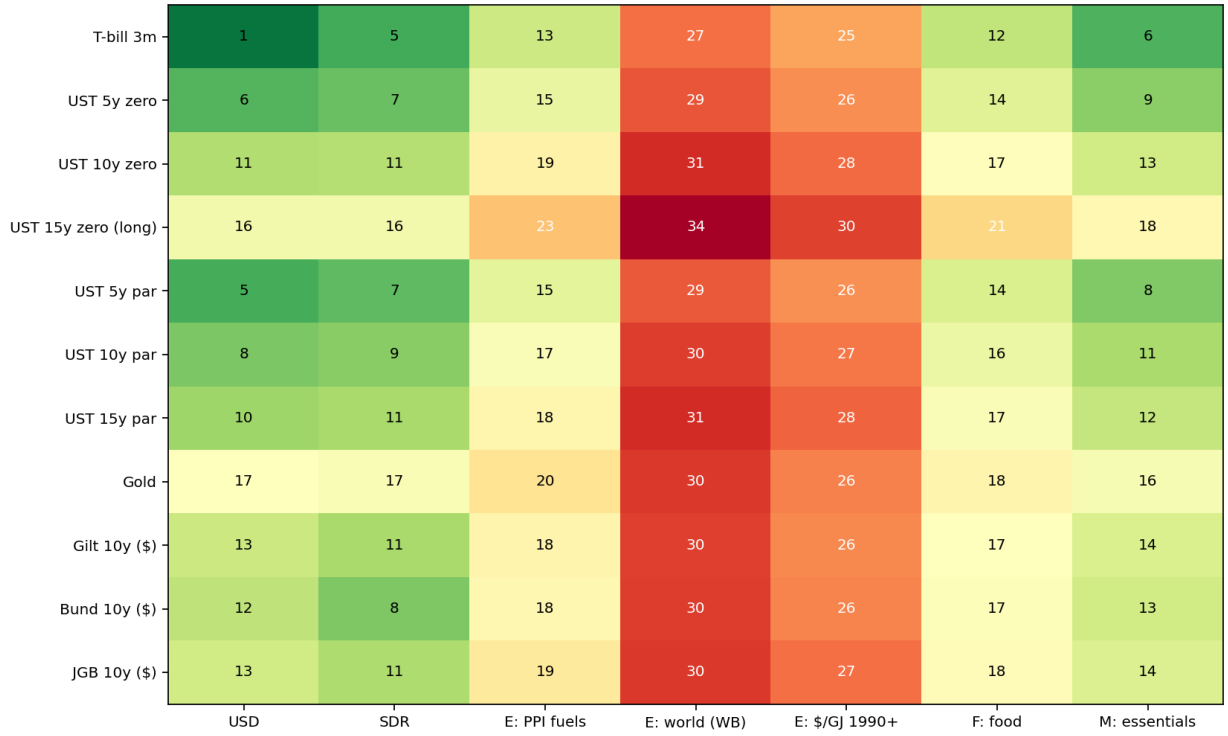


Figure 5: Annualized asset volatility by measuring unit.

C.2 Break-Even Covariance Diagnostic

A useful diagnostic for the results in Section 5.2 follows from the log-return approximation. An asset lowers liability-unit variance relative to a zero-return dollar position only if its covariance with the liability is sufficiently positive. In correlation form, the condition is

$$\rho_{i,N} > \rho^*_{i,N} = \frac{\sigma_i}{2\sigma_N}.$$

Table 9 applies this condition to the conventional reserve menu in the world-energy frame.

At the monthly horizon, no asset in the baseline menu clears its break-even correlation against the world-energy liability. Gold has the strongest positive covariance among the conventional assets, but its correlation of 0.13 remains below the approximately 0.30 required to compensate for its own volatility. Longer-duration government bonds have negative correlations with energy prices and therefore lie further below the relevant threshold.

The horizon results for gold are consistent with the term structure reported in Section 5.2. Its correlation with energy prices increases at longer horizons and exceeds the break-even value at sixty months. This provides the main source of the positive five-year hedgeable share reported in the main text.

Table 9: Break-even correlation diagnostic, world-energy frame. An asset reduces liability-unit variance relative to the zero-return dollar benchmark when $\rho > \rho^*$.

Panel A: monthly horizon					
	σ_i	ρ	ρ^*	Margin	Result
Treasury bill	1.0	+0.01	0.02	−0.01	No
UST 5y (zero)	6.2	−0.10	0.12	−0.22	No
UST 10y (zero)	11.1	−0.14	0.22	−0.36	No
UST 15y (zero)	16.0	−0.16	0.31	−0.47	No
UST 5y (par)	5.4	−0.11	0.11	−0.21	No
UST 10y (par)	8.2	−0.15	0.16	−0.31	No
UST long (par)	10.0	−0.17	0.20	−0.37	No
Gold	15.6	+0.13	0.30	−0.17	No
U.K. 10y	12.7	+0.03	0.25	−0.22	No
Germany 10y	12.1	+0.01	0.24	−0.23	No
Japan 10y	13.1	−0.06	0.26	−0.31	No
Panel B: gold across horizons					
	σ_i	ρ	ρ^*	Margin	Result
1 month	15.6	+0.13	0.30	−0.17	No
12 months	20.0	+0.22	0.32	−0.10	No
24 months	21.7	+0.35	0.36	−0.00	No
60 months	21.6	+0.68	0.42	+0.26	Yes

C.3 Full Term Structure and Measurement Diagnostics

Table 10 reports the complete set of liability specifications and the three-month timing diagnostic omitted from the main text; Figure 6 plots the world-energy and essentials term structures against their unhedged benchmarks.

Table 10: Hedging deficit and hedgeable share across liabilities and horizons.

	E (world)		E (basket)		Food		M (essentials)	
Horizon	HD	HS	HD	HS	HD	HS	HD	HS
1 month	26.1	+0.005	24.4	+0.006	11.3	+0.021	4.8	+0.001
12 months	34.0	−0.055	36.0	−0.022	15.0	−0.011	7.9	+0.020
24 months	32.1	−0.078	33.5	−0.038	15.7	−0.116	8.1	+0.117
60 months	22.7	+0.139	19.3	+0.204	14.3	−0.410	6.5	+0.274
3 months (diagnostic)	34.0	−0.010	32.2	−0.004	14.6	+0.009	6.3	−0.014

The three-month specification addresses the timing difference between end-of-month financial prices and goods-price indices constructed from monthly averages. The world-energy deficit rises from 26.1 percent in the monthly specification to 34.0 percent on the three-month diagnostic. The direction of the difference indicates that the monthly timing convention attenuates rather than amplifies the estimated goods-price exposure.

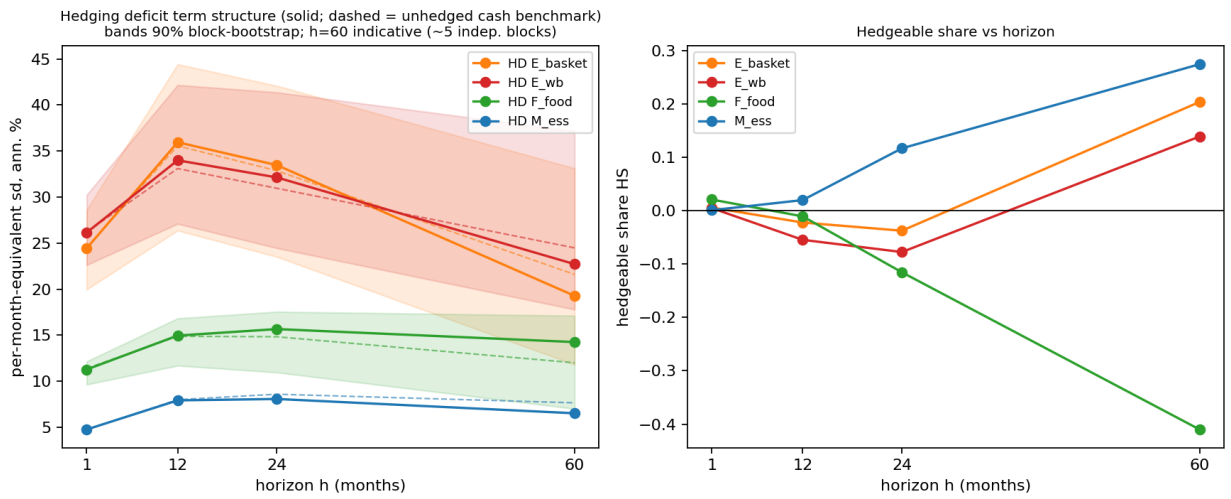


Figure 6: Term structure of the hedging deficit: per-month-equivalent HD (solid) against the unhedged benchmark (dashed) at the pre-registered horizons, with 90 percent moving-block bootstrap intervals.

Alternative zero-coupon and par-coupon Treasury constructions produce very similar hedging deficits in the goods frames. Alternative block lengths for the moving-block bootstrap likewise have little effect on the reported point estimates.

The food results are less stable across long horizons than the energy results. In particular, the five-year full-sample hedgeable share becomes negative. Because this pattern is concentrated in long windows spanning different real-interest-rate and food-price regimes, the food term structure is treated as an additional specification rather than a headline result.

C.4 Import-Cover Distributions and Welfare Comparison

Figure 7 reports the complete twelve-month import-cover distributions underlying Table 4.

The similarity in fifth-percentile outcomes is not limited to the reported quantile. The broader distributions of energy-import cover overlap substantially across portfolios optimized under different measuring units. Differences are larger in dollar volatility than in liability-unit downside outcomes.

The secondary welfare comparison yields the same qualitative conclusion. At $\gamma = 5$, the annual certainty-equivalent loss from selecting the dollar-optimal portfolio when performance is evaluated in world-energy units is approximately 17–29 basis points under the two expected-return specifications described in Section 4. The corresponding estimate for the essentials basket is approximately 2–3 basis points.

These quantities should not be interpreted as estimates of the welfare cost of the full goods-price exposure. Cross-numeraire regret compares two feasible portfolios within the same conventional opportunity set. When that set contains little covariance with the liability, the difference between its internally optimal portfolios can be small even when the common residual exposure is large.

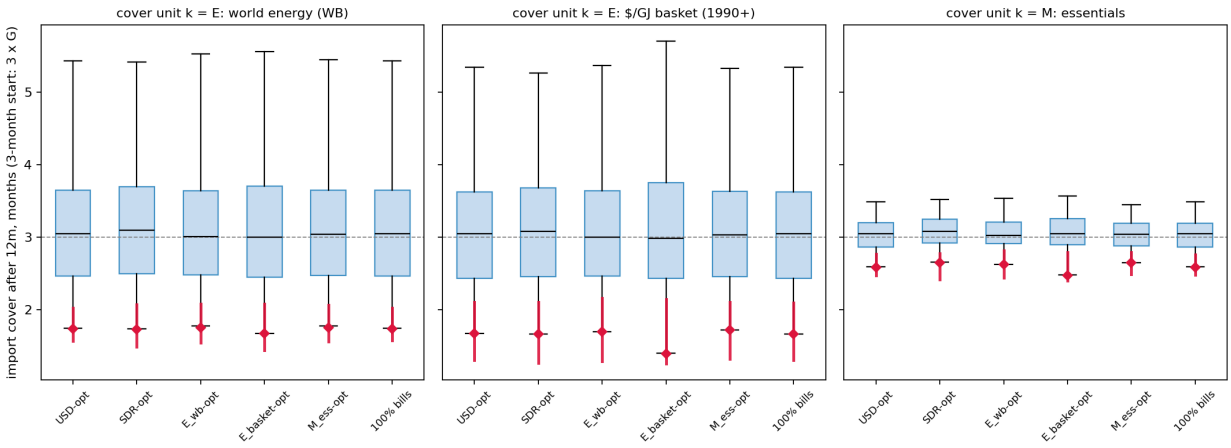


Figure 7: Out-of-sample twelve-month import-cover distributions by portfolio mandate. Portfolios are formed on rolling 120-month windows and evaluated over the subsequent year.

C.5 Conditional Risk Structure

Figure 8 reports the conditional risk structure underlying Section 5.4: annualized volatility by state and measuring unit on the post-1983 identification sample, together with the energy-frame correlation matrices estimated within the supply state and within its complement. Realized energy-frame volatility in supply months is somewhat lower than in the complement, and the compressed correlation structure documented in Figure 2 is present in both states.

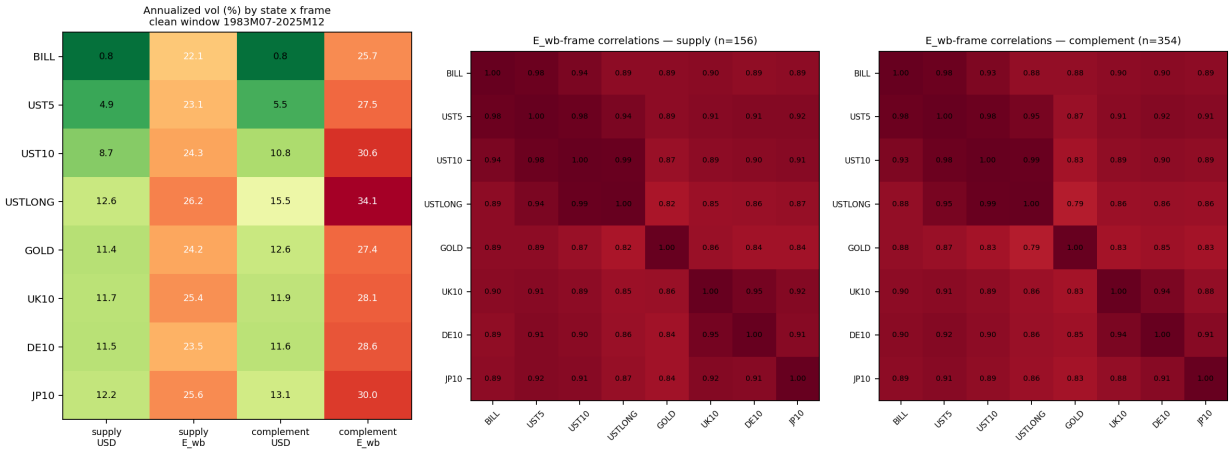


Figure 8: Conditional risk structure, post-1983 identification sample. Left: annualized volatility by state and measuring unit. Center and right: energy-frame correlation matrices within the supply state and the complement.

C.6 Event Exercises

The event exercises provide an ex-ante complement to the conditional covariance estimates. Portfolio weights are estimated using information available before each event and are then

held fixed over the event window.

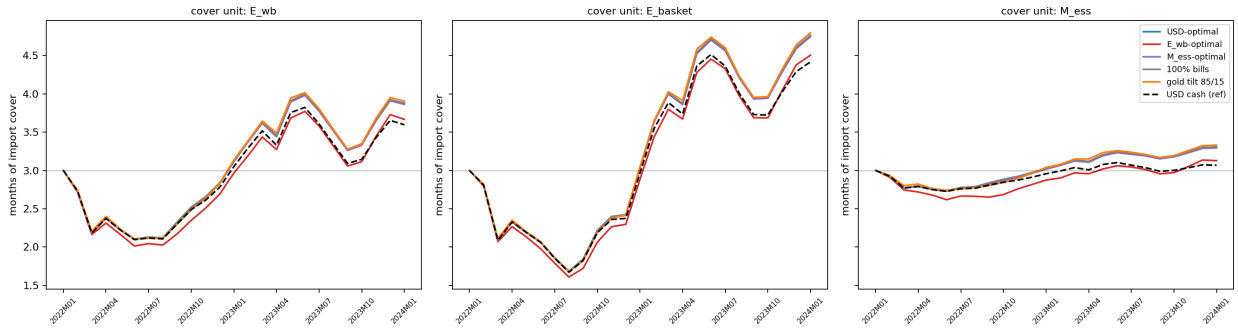


Figure 9: Energy-import cover during the 2022 episode using portfolio weights formed with information available through the end of 2021. All paths begin at three months of cover.

During the 2022 episode (Figure 9), a bills-only portfolio declines from three months of energy-import cover to approximately 2.1 months by mid-year. The portfolio optimized ex ante in energy units reaches a slightly lower trough of approximately 2.0 months because its pre-event allocation included longer-duration bonds, which also experienced losses as interest rates rose. A portfolio containing a moderate gold allocation produces a path close to the bills-only benchmark around the trough.

The exercise does not imply that energy-unit optimization systematically worsens reserve performance. Rather, it illustrates the limited degree of protection available from the conventional asset set when the joint realization of asset returns and energy prices differs from the historical covariance structure used at the formation date.

The 2026 exercise (Figure 10) is treated separately because the observations are preliminary and are not part of the estimation sample.

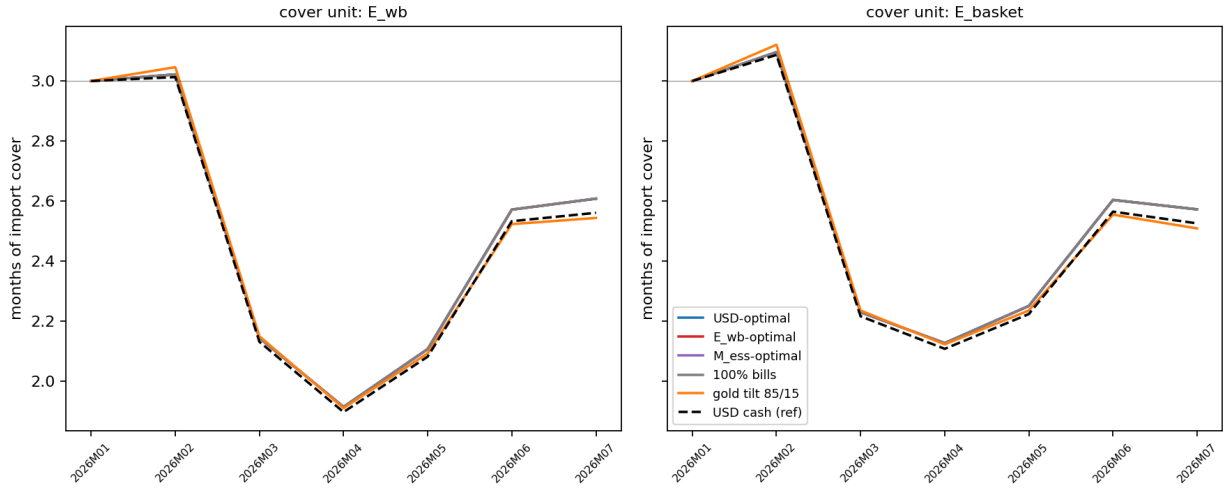


Figure 10: Preliminary 2026 out-of-sample energy-import-cover paths. Portfolio weights are formed using pre-event information.

By the beginning of the 2026 episode, the alternative minimum-risk portfolios had converged to allocations dominated by Treasury bills. Their subsequent import-cover paths are therefore nearly identical. From the standardized three-month starting value, energy cover falls to approximately 1.9 months by April before partially recovering later in the available sample. These estimates remain preliminary because the final 2026 data vintage is not yet available.

D Additional Results for the Expanded Hedging Set

This appendix reports the additional results underlying Section 6. It contains the complete financial-instrument staircase, spanning tests, futures implementation diagnostics, constrained-sleeve specifications, state-dependent results, and additional information on physical strategic stocks and storage costs.

D.1 Financial-Instrument Staircase

Table 11 reports the cumulative effect of adding each financial instrument to the conventional reserve menu on the common 1999M02–2024M04 sample.

Table 11: Expansion of the financial opportunity set, world-energy liability. HS denotes the cumulative hedgeable share. HK reports the p-value from the Huberman–Kandel spanning test of each addition relative to the preceding opportunity set. Commodity-futures positions use the five-business-day roll convention. The shrinkage estimator is re-estimated for each opportunity set, so the cumulative share can decline by trivial amounts (at most about 0.002) when an added asset receives no weight; under a fixed covariance matrix the sequence is monotone by construction. Increments are computed from unrounded values.

Asset set	HS, $h = 1$	ΔHS	HK p	HS, $h = 12$
Conventional reserve menu	0.018	—	—	0.040
+ WTI futures	0.517	+0.499	10^{-43}	0.815
+ Henry Hub futures	0.524	+0.007	0.049	0.822
+ Energy equities	0.522	−0.001	0.003	0.822
+ Inflation-linked Treasuries	0.523	+0.000	0.0005	0.821
+ Energy-indexed bill	0.910	+0.387	≈ 0	0.908

The WTI futures position accounts for nearly all of the improvement produced by the traded financial instruments. At the monthly horizon, the optimizer assigns approximately 53 percent of the portfolio to the collateralized WTI exposure. At twelve months, the allocation reaches 70 percent, the maximum permitted by the 30 percent Treasury-bill floor.

Henry Hub futures provide substantial hedging capacity when considered alone at longer horizons (approximately 0.53 at twelve months, versus approximately 0.05 at one month), but add relatively little after crude exposure has already entered the portfolio. This reflects the large common component in oil and gas prices as well as the relatively smaller gas weight in the world-energy index.

Energy equities and inflation-linked Treasuries provide little incremental reduction in energy-liability variance. The formal spanning tests for these assets can reject because the tests consider both mean and variance dimensions. The minimum-variance exercise uses only the latter. For example, adding TIPS on a standalone basis leaves the twelve-month hedging deficit essentially unchanged (a change of approximately 0.05 percentage points, bootstrap p-value of approximately 0.9). The small declines in the cumulative share visible in Table 11 after these additions (at most about 0.002) are an artifact of re-estimating the shrinkage covariance matrix for each opportunity set: with the covariance matrix held fixed, adding an asset cannot worsen the constrained minimum, and the optimizer in fact assigns these assets zero or near-zero weight.

In the broader essentials specification, the ranking is somewhat different. Energy equities provide a larger contribution at the annual horizon, where an approximately 18 percent allocation raises the hedgeable share from roughly 0.63 to 0.75. The result is consistent with the smoother and more diversified exposure represented by the broader import basket.

D.2 Futures Returns and Roll Conventions

All commodity-futures positions in the baseline extension are fully collateralized. The return therefore equals the Treasury-bill return plus the excess return generated by maintaining and rolling the futures position.

The realized roll component varies considerably through time. For WTI, roll returns were strongly positive during the backwardated 1980s and approximately neutral to moderately positive in the 1990s, before becoming negative during the more persistent contango episodes of the 2000s and 2010s. Approximate annualized roll contributions by decade are:

1980s	+20%
1990s	+2%
2000s	−5%
2010s	−9%

These quantities are not additional charges imposed after the portfolio calculation; they are already embedded in the futures return used in the optimization.

Henry Hub futures have historically exhibited larger negative roll costs, ranging from approximately −18 to −34 percent per year across the decades in which the contract is available under the headline roll convention. This helps explain why natural-gas futures provide substantially less incremental value once crude futures are present.

Two roll implementations are used as a robustness check. Their return series are correlated at approximately 0.99 for WTI and 0.98 for Henry Hub, and changing the convention alters cumulative hedgeable shares by less than approximately 0.01 in the energy frame at the monthly horizon (under 0.002 at twelve months) and by less than approximately 0.02 in the essentials frame.

D.3 Basis Risk

The financial hedge does not replicate the world-energy liability exactly. The principal WTI contract references a particular crude benchmark, while the liability measure contains a

broader set of internationally traded energy prices.

Over the common sample, the WTI hedge and world-energy liability have a correlation of approximately 0.90 and annualized tracking error of approximately 14 percent. Tracking error is smaller relative to the crude component alone and larger relative to the broader dollars-per-gigajoule energy basket containing crude oil, natural gas, and coal.

The basis is also regime-dependent. WTI traded at an average premium to Brent before the early 2010s and at a discount thereafter. Consequently, even a financial hedge with high unconditional covariance can experience changes in effectiveness when the relationship among regional energy benchmarks shifts.

The distinction between the traded and frictionless claims can therefore be interpreted as a decomposition of the remaining financial hedging deficit. The energy-indexed bill removes the basis between the contract and the liability and does not require repeated futures rolls. The residual exposure under that benchmark is primarily generated by the maintained Treasury-bill liquidity floor.

D.4 Constrained Commodity Sleeves

The baseline extension allows the optimizer to determine the size of the commodity exposure subject only to the general portfolio constraints. To evaluate a more limited implementation, we also cap the combined non-baseline allocation.

With a 10 percent maximum commodity-linked sleeve, the world-energy hedgeable share is approximately 0.19 at the monthly horizon and 0.23 at twelve months. The sleeve constraint binds. Both quantities are therefore substantially below the 0.52 and 0.82 obtained when the optimizer is allowed to select the exposure endogenously.

The comparison is useful for relating the exercise to reserve-allocation proposals that contemplate relatively small commodity positions. Such sleeves move the portfolio in the direction implied by the liability exposure but do not span a large share of the liability when the desired hedge is substantially larger than the permitted position.

The difference is smaller for the broader essentials basket. Because energy constitutes only one component of that liability, the unconstrained optimizer itself demands a more modest commodity allocation at the monthly horizon. At twelve months the desired allocation is substantially larger, and the 10 percent cap binds for the essentials liability as well.

D.5 State-Dependent Performance of the Financial Extension

The financial extension remains effective when the sample is restricted to adverse supply states, but its estimated effectiveness is lower than in the unconditional sample.

On the common sample for which both the state indicator and futures data are available, the conventional reserve menu produces a hedgeable share of approximately 0.03. Adding the WTI-based exposure raises the conditional hedgeable share to approximately 0.28. The 90 percent bootstrap interval is approximately $[0.16, 0.46]$.

This represents a substantial improvement relative to the conventional reserve set but approximately half the unconditional variance reduction. One interpretation is that the basis between WTI and the broader world-energy liability widens during supply disruptions. The same state in which direct energy exposure is particularly valuable can therefore also be

a state in which a single traded benchmark provides a less complete representation of the import liability.

The horizon comparison points in the opposite direction. At twelve months, the financial overlay performs more strongly than at the monthly horizon. The baseline world-energy hedging deficit rises substantially between one and twelve months, whereas the expanded portfolio reduces much of this medium-horizon persistence. The result is consistent with persistent changes in energy prices being more easily hedged than high-frequency basis and measurement noise.

D.6 Strategic Physical Stocks

Table 12 summarizes selected strategic stocks relevant to the physical extension. These quantities are included to document the institutional scale on which physical inventories are already maintained. They are not inputs into the baseline portfolio estimates.

Table 12: Selected sovereign strategic inventories. Quantities are those used in the current data vintage; entries based on non-official estimates are explicitly identified. Primary sources and date-specific documentation are recorded in the replication files.

Holder	Instrument	Approximate quantity / requirement	Status
United States	Strategic Petroleum Reserve	355 million barrels, May 2026	Official
IEA members	Emergency oil stocks	Minimum 90 days of net imports under IEA obligation	Official
Japan	State and mandated private petroleum stocks	Approximately 470 million barrels at end-2025	Official
Germany	Petroleum stocks	Approximately 110 million barrels crude plus 67 million barrels products	Official
China	Government crude stocks	Several hundred million barrels	Estimate
India	Strategic petroleum caverns	Approximately 39 million barrels capacity	Official
India	Central grain stocks	Approximately 60 million tonnes, April 2026	Official
China	State grain stocks	Large but not officially disclosed	Estimate

The table illustrates two features relevant to the main text. First, strategic commodity inventories are already maintained at quantities that are economically large relative to import needs. Second, the institutional form varies substantially across countries, including direct government ownership, statutory stockholding obligations, and mandated private inventories.

Because reporting standards differ and some major holders do not publish complete inventory data, the table should not be interpreted as a cross-country ranking of strategic-stock adequacy.

D.7 Storage Costs and the Physical-Hedge Calculation

The physical extension models stored crude as a position whose value changes with the spot price of crude net of a storage charge. Several storage-cost assumptions are considered.

Commercial crude storage has historically been priced at approximately \$0.25–\$0.60 per barrel per month in relevant storage markets. The operating cost implied by the U.S. Strategic Petroleum Reserve budget is substantially lower, approximately \$0.04–\$0.05 per barrel per month in the current calibration. These measures are not directly comparable in every respect: commercial tank rates include market scarcity rents and capital costs, whereas government operating budgets reflect previously constructed infrastructure.

The covariance-based hedgeable-share result is relatively insensitive to the assumed carrying cost because the storage charge primarily affects expected returns rather than the contemporaneous co-movement between the stored commodity and the liability. At the twelve-month horizon, stored crude produces a hedgeable share of approximately 0.88, compared with 0.82 for the futures-based financial extension and approximately 0.91 for the frictionless liability-linked benchmark.

The remaining gap arises largely because crude does not span the natural-gas and coal components of the world-energy index.

D.8 Joint Financial and Physical Cover

The joint-cover exercise begins with a sovereign position equal to three months of energy imports. A fraction s is represented by stored crude and $1 - s$ by the financial reserve position. The physical share is restricted to

$$0 \leq s \leq 0.5,$$

which is intended to represent a capacity constraint rather than an estimated optimum.

Under sovereign-cavern storage costs, the fifth-percentile twelve-month cover outcome increases approximately monotonically with s. The principal endpoints are:

Configuration	5th-percentile cover	Approx. expected-cover cost
All financial	1.73	—
Half physical, sovereign storage cost	2.30	0.125 months
Half physical, commercial storage cost	2.11	0.33 months

When the evaluation is restricted to adverse supply-state windows, the lower-tail outcome rises from approximately 1.62 months in the all-financial case to approximately 2.28 months at the half-physical allocation under the sovereign-cost assumption.

These estimates are illustrative rather than optimal-policy calculations. They do not incorporate endogenous storage construction, release timing, replenishment decisions, product substitution, or general-equilibrium effects.

D.9 Interpretation and Scope of the Physical Extension

The physical-stock calculation captures only part of the economic value of strategic inventories. Valuing the stock through its commodity-price return treats it as price insurance for the sovereign holding the inventory. Strategic releases can additionally provide quantity insurance by relaxing a physical supply constraint. To the extent that releases affect market

availability or prices for other users, they can also have public-good effects that lie outside the portfolio framework.

The physical hedge is also commodity-specific. Stored crude does not eliminate exposure to refined-product spreads, natural-gas prices, coal prices, or infrastructure disruptions. Its effectiveness therefore depends on the correspondence between the stored asset and the actual crisis-period import liability.

Finally, physical inventories and foreign-exchange reserves serve different liquidity functions. Financial reserves can be used immediately for intervention, debt service, and a broad range of external payments. Strategic commodities must be released, transported, refined where necessary, or sold before they can serve those functions. The physical extension should therefore be interpreted as an expansion of the sovereign hedging set rather than an expansion of the conventional central-bank reserve portfolio.

E Additional Robustness Results

This appendix reports additional specifications underlying Section 7. The exercises are grouped by the component of the empirical design they modify: liability measurement, asset construction, portfolio restrictions, estimation, sample choice, state identification, and reserve-adequacy calculations. Unless otherwise noted, each specification changes one element of the baseline design while retaining the remaining assumptions.

E.1 Alternative Energy and Import Liabilities

The headline energy specification uses the World Bank world energy price index. We consider several alternatives intended to determine whether the low hedgeable share is specific to that index.

Alternative energy measures include the constructed internationally traded energy basket under current consumption weights, weights fixed at the beginning of the modern subsample, and an oil-only specification. Across these alternatives, the monthly hedgeable share remains between approximately 0.005 and 0.006. The level of the hedging deficit moves with the volatility of the underlying price series, but the variance reduction attributable to the conventional reserve portfolio remains close to zero.

For the broader essentials measure, replacing the interpolated manufactures-price component with the available monthly non-fuel import-price series changes the hedging deficit by approximately 0.1 percentage point on the common sample. Alternative weighting conventions likewise change the volatility of the liability more than they change the portfolio's ability to hedge it.

The food specification is less stable at long horizons. In the full sample, the five-year hedgeable share becomes substantially negative. Splitting the sample around major changes in real-interest-rate and food-price regimes produces positive five-year hedgeable shares within the principal subsamples. The full-sample result therefore appears to reflect long-horizon interactions between compounded short rates and secular food-price movements rather than a robust inability to hedge food at all horizons. Because energy is the principal

stress liability in the paper, this long-horizon food result is treated as a secondary specification.

E.2 National Import-Price Measure

Türkiye provides the principal country-specific check because a monthly national import unit-value index is available over a sufficiently long recent sample and because earlier reserve-composition work found that the choice of import-price deflator materially changed its estimated currency allocation.

Using the published Turkish import unit-value index as the liability produces a hedgeable share of approximately -0.01. The bills-only twelve-month IMaR is approximately 2.3 months from the standardized three-month initial position.

The national index correlates at approximately 0.47 with the essentials basket, 0.43 with the India-weighted basket, and 0.40 with the world energy index (0.31 with the food index). The exercise therefore does not simply reproduce the common index under a country-specific label. Despite the difference in measured price dynamics, the conventional asset menu continues to provide little variance reduction.

The principal limitation is sample length. The currently available national series begins in 2013 on the public statistical portal and therefore cannot replicate the long historical analysis conducted with the common price indices.

E.3 Alternative Government-Bond Return Construction

The baseline exhibits use constant-maturity zero-coupon returns for U.S. Treasury securities where available. We reconstruct the corresponding analysis using par-coupon returns from the same underlying government yield curves.

The alternative convention materially changes the dollar volatility of longer-duration bonds but has very little effect on the goods-frame portfolio statistics. Across the principal goods liabilities, replacing the zero-coupon Treasury legs with their par-coupon counterparts changes the hedging deficit by no more than approximately 0.03 percentage points.

This result is consistent with Section 5.1. Once the common goods-price component accounts for a large share of transformed-return variance, differences in bond-duration construction become relatively small for the liability-unit portfolio problem.

E.4 Timing of Financial and Goods Prices

Several goods-price series are reported as monthly averages, whereas most financial returns are constructed using end-of-month observations. This difference can attenuate measured contemporaneous covariance.

The three-month diagnostic described in Section 4 reduces the importance of this timing mismatch. For the world-energy liability, the per-month-equivalent hedging deficit rises from approximately 26.1 percent in the monthly specification to approximately 34.0 percent in the three-month diagnostic. The corresponding hedgeable share remains approximately zero.

The direction of the difference implies that the baseline monthly calculation is not generating the hedging deficit through timing mismatch. If anything, the mixed observation convention reduces the measured liability exposure.

E.5 Currency-Hedged Foreign Government Bonds

The baseline treats foreign government bonds as unhedged because reserve currency composition is part of the allocation decision. As an alternative, we hedge the foreign-currency exposure of the German, Japanese, and U.K. government bonds back into dollars.

Currency hedging reduces the dollar volatility of these instruments substantially. On the common sample used for the comparison, however, the world-energy hedging deficit is approximately 29.7 percent with the baseline unhedged bonds and 29.9 percent with the dollar-hedged versions. The difference is statistically and economically small.

The same transformation has a large effect in the SDR exercise. The SDR hedgeable share declines from approximately 0.73 to 0.17 when foreign bond returns are hedged into dollars because the hedge removes the foreign-currency exposure that allows the portfolio to replicate components of the SDR basket.

The contrast illustrates that currency hedging is effective when currency exposure is the liability being managed. It does not generate covariance with a goods-price liability that was absent from the underlying assets.

E.6 Portfolio Constraints

We consider several alternatives to the baseline long-only, fully invested portfolio with a 30 percent Treasury-bill floor.

Imposing a 25 percent maximum exposure to each non-U.S. issuer has little effect on the goods-frame results. In the baseline minimum-risk portfolios, the Treasury-bill floor is not binding against the principal goods liabilities because the optimizer voluntarily holds substantially more than the required allocation.

The active restrictions are instead the non-negativity constraints on selected duration exposures. In an unconstrained diagnostic, the optimizer attempts to obtain positive exposure to energy by taking negative positions in government bonds whose historical covariance with energy is negative. Allowing these positions improves the hedging deficit only modestly and produces allocations that are not representative of conventional reserve management. The diagnostic is therefore informative about the sign of the covariance structure rather than an alternative reserve portfolio.

E.7 Covariance Estimation and Bootstrap Design

The baseline covariance matrices use Ledoit–Wolf shrinkage. Re-estimating the world-energy portfolio using the conventional sample covariance matrix produces an annualized hedging deficit of approximately 26.11 percent, compared with approximately 26.15 percent under shrinkage.

Changing the moving-block bootstrap length from the baseline 24 months to 12 or 36 months changes the endpoints of the reported confidence intervals by at most approximately 0.4 percentage point for the headline hedging deficit.

These alternatives affect estimated uncertainty more than the central point estimates.

E.8 Subperiod Stability

To assess the influence of particular historical regimes, the full sample is divided into four broad subperiods, split at 1985, 1999, and 2013; the first subperiod contains the thin-market early-JGB segment discussed in Appendix B. The estimated world-energy hedgeable shares are approximately

$$0.02, \qquad 0.09, \qquad 0.07, \qquad 0.00.$$

The composition of the estimated minimum-risk portfolio varies substantially across these periods. Gold's estimated optimal share, for example, ranges from zero to approximately 62 percent.

The instability of portfolio weights does not translate into a subperiod in which conventional reserve assets span a large share of the energy liability. The exercise therefore suggests that the low hedgeable share is not driven exclusively by the major oil shocks of the 1970s or by a single monetary regime.

E.9 Alternative Supply-State Definitions

Table 13 reports the complete state-identification battery.

Table 13: Alternative adverse supply-state definitions. HD is per-month-equivalent annualized percent.

Supply-state definition	n	HD	HS
Signed news, 0.5σ threshold	156	21.8	0.016
Signed news, 1.0σ threshold	79	25.9	0.020
Unsigned, $\lvert z\rvert > 0.5\sigma$	307	33.0	0.019
Structural taxonomy	112	20.7	0.014
Ready-style implementation	123	16.7	0.022
Announcement-only	48	29.6	0.036
Pre-1983 extrapolated observations	*30*	*13.5*	*0.073*

The first six rows use classifications that can be implemented on the instrument-supported sample. Their hedgeable shares lie between 0.014 and 0.036 despite substantial differences in the volatility of the observations assigned to the supply state.

The final row uses pre-1983 values for which the supply-news measure requires model-based extrapolation. Its hedgeable share is higher, approximately 0.073, but the sample consists of only 30 observations and is excluded from all headline conditional results.

The pre-pandemic vintage of the supply-news data produces 89–92 percent agreement in state membership relative to the later vintage under the principal definitions. Re-estimating the conditional portfolio with that vintage changes the hedging deficit by approximately 0.05 percentage point.

E.10 Import-Months-at-Risk Without Historical Drift

The baseline IMaR exercise resamples observed joint return histories and therefore retains the sample's historical average returns. To assess whether the lower-tail result is driven by drift rather than covariance and tail realizations, we repeat the exercise after demeaning the underlying return paths.

The fifth-percentile world-energy cover outcome moves from approximately 1.73 months in the baseline exercise to approximately 1.63 months in the demeaned variant. The change is modest relative to the overall reduction from the standardized three-month starting position.

The result indicates that historical drift is not responsible for the estimated downside purchasing-power exposure.

E.11 Break-Even Census and Sample Convention

The break-even diagnostic reported with the Section 5 results depends modestly on the sample convention used to estimate individual asset correlations.

On the joint estimation window used for the portfolio exercises, one of the 88 asset-liability-horizon cells clears its break-even correlation: gold against energy at the five-year horizon. When each pair is instead evaluated over its longest available sample, seven cells pass. The additional cases are driven mainly by gold's 1971–74 revaluation period and by the Treasury bill narrowly clearing its very small break-even threshold in several specifications.

The monthly conclusion is essentially unchanged across the two registers: conventional reserve assets generally fail to provide sufficient positive covariance with the goods liability to reduce liability-unit variance materially.

E.12 Additional Data Limitations

Four additional data limitations warrant note.

First, a separate historical bank-deposit leg is not included because the relevant long eurodollar benchmark was discontinued and a consistent freely redistributable replacement is not available over the full sample. The Treasury-bill leg provides the corresponding short-duration dollar exposure in the baseline portfolio.

Second, constructed government-bond returns can be compared with commercial total-return indices over portions of the sample, but the complete licensed histories cannot be included in the public replication package.

Third, the Türkiye national-price exercise is restricted to the period made available through the current public statistical portal.

Fourth, 2026 observations remain preliminary and are excluded from the baseline estimation sample.

These limitations affect particular extensions and validation exercises rather than the central comparison between conventional reserve assets and the principal goods-price liabilities.

E.13 Sanctions and Asset Usability

The analysis in Section 7.3 treats sanctions and asset-access restrictions as conceptually separate from the baseline goods-price exposure. A simple extension of the framework can make the distinction explicit.

Let $V_{i,t}$ denote the market value of reserve asset i, and let $a_{i,s} \in [0,1]$ denote its usable fraction in state s. Effective reserve value is then

$$\widetilde{V}_{i,t,s} = a_{i,s} V_{i,t}.$$

The liability-unit value of the asset becomes

$$\frac{\widetilde{V}_{i,t,s}}{P_{N,t}} = a_{i,s} \frac{V_{i,t}}{P_{N,t}}.$$

Goods-price risk operates through variation in $P_{N,t}$ relative to asset value. Usability risk operates through $a_{i,s}$. The two can occur separately or jointly.

This decomposition also clarifies the role of custody. An asset may have a favorable $a_{i,s}$ under one sanctions configuration but face lower market liquidity or higher mobilization costs under another custody arrangement. These features would enter a richer reserve model as separate state-dependent constraints or haircuts rather than as changes in the asset's covariance with the import liability.

The paper does not estimate $a_{i,s}$, model strategic sanctions behavior, or optimize reserve holdings jointly over goods-price and sanctions states. The purpose of the extension is solely to distinguish the purchasing-power risk studied in the main analysis from a separate and economically important dimension of reserve safety.

F Reserve Objectives, Risk Measures, and Institutional Scope

This appendix clarifies the relationship among the reserve-management functions discussed in the paper. It is intended to distinguish the import-purchasing-power exposure studied in the empirical analysis from other objectives that can imply different measuring units, asset characteristics, and institutional arrangements.

F.1 Reserve Functions and Relevant Risk Objects

International reserves serve several purposes, and the economic object against which their performance is evaluated need not be the same across those purposes. Table 14 summarizes the principal distinctions used in the paper.

Table 14: Reserve functions and associated risk objects. The entries describe analytically distinct dimensions of reserve management rather than alternative prescriptions for portfolio construction.

Reserve function	Relevant unit or state variable	Principal exposure	Connection to this paper
Foreign-exchange intervention	Intervention currency or currency basket	Exchange-rate risk, market liquidity	Monetary reserve function
External debt service	Currency and maturity of external liabilities	Currency mismatch, refinancing and liquidity risk	Monetary liability matching
Import financing	Import prices or relevant goods basket	Relative-price and purchasing-power risk	Baseline liability-unit analysis
Reserve usability	State-contingent accessible asset value	Sanctions, custody, settlement, convertibility	Section 7 additional consideration
Physical supply continuity	Quantity of the relevant essential good	Availability, logistics, product and infrastructure risk	Section 6 physical extension

The categories are not mutually exclusive. A single reserve asset can simultaneously provide intervention liquidity, satisfy a debt-service requirement, and retain some purchasing power over imports. The purpose of the decomposition is instead to make clear that performance along one dimension does not establish performance along another.

This point is especially relevant to the distinction between currency matching and goods-price hedging. Matching the currency composition of reserves to the invoice currency of imports reduces one component of the liability. The remaining movement in the price of the imported good is a separate relative-price exposure. The transformation in Equation (1) isolates that second component.

F.2 Stock Adequacy and Purchasing-Power Adequacy

Let initial reserve cover equal C_t months of imports. A conventional import-cover statistic evaluates this quantity at time t. Equation (5) in the main text shows that future cover depends on the joint evolution of the reserve portfolio and the relevant import price:

$$\frac{C_{t+H}}{C_t} = \prod_{\tau=1}^{H} \frac{1 + R^{\$}_{p,t+\tau}}{1 + R^{\$}_{N,t+\tau}}.$$

The initial level C_t and the distribution of C_{t+H}/C_t therefore contain different information. The former describes the size of the initial financial buffer. The latter additionally reflects uncertainty about the value of that buffer relative to the liability.

The distinction does not require that import purchasing power become the objective for the entire reserve stock. It applies to the component of reserve adequacy for which the quantity of external goods financeable by the reserve position is economically relevant.

The empirical results in Section 5 illustrate this separation. Monetary volatility varies considerably across alternative conventional portfolio mandates, whereas their lower-tail outcomes in energy-import units are similar. The difference reflects the large common liability component that remains after optimization within the conventional opportunity set.

F.3 Interpretation of the Expanded Opportunity Set

The extensions in Section 6 can be interpreted as progressively changing the relationship between the payoff of the available asset and the import liability.

Table 15: Interpretation of alternative instruments in the expanded opportunity set.

Instrument	Connection to import liability	Principal residual exposures
Conventional reserve assets	Primarily monetary claims; limited direct covariance with energy prices	Goods-price risk remains largely unspanned
Commodity futures	Direct exposure to a traded commodity benchmark	Basis risk, rollover returns, collateral and position constraints
Hypothetical liability-indexed bill	Direct exposure to the liability index by construction	Maintained reserve constraints; hypothetical market availability
Stored crude	Direct ownership of one component of the physical liability	Storage cost, product basis, infrastructure and release constraints

The ordering in Table 15 is analytical rather than normative. Moving down the table increases the directness of exposure to the energy liability in the empirical exercises, but it also changes the liquidity, operational, and institutional characteristics of the instrument.

The hypothetical indexed bill is useful because it separates two sources of incomplete hedging. The difference between the conventional reserve menu and the indexed benchmark reflects the absence of direct liability exposure in the baseline asset set. The difference between traded commodity futures and the indexed benchmark reflects, in part, basis and implementation frictions associated with obtaining similar exposure through existing financial markets.

Physical inventories require an additional distinction. Their market value can be incorporated into the same relative-price framework, but their economic role is not exhausted by that valuation. Ownership of the commodity can also provide physical availability in states in which procurement is constrained. The portfolio calculations in Section 6 capture the price component and do not attempt to value this additional quantity-insurance channel.

F.4 Financial and Physical Balance Sheets

The financial and physical instruments considered in the paper are generally held by different public institutions and governed by different mandates. Foreign-exchange reserves are commonly managed by central banks or monetary authorities, whereas strategic commodity inventories may be held by energy agencies, finance ministries, dedicated stockholding agencies, or mandated private entities.

The empirical analysis does not treat these arrangements as institutionally equivalent. Rather, it identifies an economic overlap: both types of assets can affect the sovereign's capacity to maintain essential imports during an external disturbance.

Conventional reserve statistics generally measure financial claims that satisfy established definitions of international reserves. Strategic commodity inventories do not satisfy those definitions and are not treated as international reserves in this paper. Their inclusion in the Section 6 extension instead changes the unit of analysis from the central-bank reserve portfolio to the broader set of sovereign instruments capable of affecting the import-capacity outcome.

F.5 Partial-Equilibrium Interpretation

The expanded-asset exercises take asset prices, market depth, futures-market capacity, and storage infrastructure as given. They therefore describe the opportunity set faced by an individual institution or sovereign at observed market prices.

The results do not characterize an equilibrium in which many reserve managers simultaneously alter their demand for commodity-linked claims. At sufficiently large scale, changes in official-sector demand could affect futures prices, basis relationships, collateral requirements, liquidity, and the relative price of conventional safe assets. Similarly, expansion of physical strategic reserves would involve endogenous storage investment and potentially affect commodity inventories and prices.

The hedgeable-share improvements reported in Section 6 should therefore be interpreted as partial-equilibrium measures of how the attainable portfolio set changes when additional instruments become available, rather than estimates of the aggregate effect of a broad change in reserve-management practice.

F.6 Scope of the Interpretation

The framework distinguishes several concepts that are sometimes combined in discussions of reserve safety:

$$\text{market value} \neq \text{liability purchasing power} \neq \text{state-contingent usability} \neq \text{physical availability}.$$

Each can be relevant to external resilience, but each is governed by a different source of risk. The empirical contribution of the paper concerns the second object: the purchasing power of financial reserve assets relative to import-price liabilities. Sections 6 and 7 show how that object relates to, but remains distinct from, direct commodity exposure, physical inventories, and reserve usability.

Nothing in this decomposition identifies an optimal institutional arrangement or portfolio allocation. It provides a common accounting of the different risks represented in the empirical exercises.

References

Adelman, M. A. (1990). The first oil price explosion 1971–1974. Working Paper MIT-CEPR 90-013WP, MIT Center for Energy Policy Research. Table VI carries the 1973–74 posted-price series.

Aizenman, J., Edwards, S., and Riera-Crichton, D. (2012). Adjustment patterns to commodity terms of trade shocks: The role of exchange rate and international reserves policies. *Journal of International Money and Finance*, 31(8):1990–2016.

Aizenman, J. and Lee, J. (2007). International reserves: Precautionary versus mercantilist views, theory and evidence. *Open Economies Review*, 18(2):191–214.

Aizenman, J. and Riera-Crichton, D. (2008). Real exchange rate and international reserves in an era of growing financial and trade integration. *Review of Economics and Statistics*, 90(4):812–815.

Arslanalp, S., Eichengreen, B., and Simpson-Bell, C. (2023). Gold as international reserves: A barbarous relic no more? *Journal of International Economics*, 145:103822.

Bank of Ghana (2026). Annual report and financial statements 2025. Technical report, Bank of Ghana.

Barsky, R. B. and Kilian, L. (2002). Do we really know that oil caused the great stagflation? A monetary alternative. In *NBER Macroeconomics Annual 2001*, volume 16, pages 137–183. MIT Press.

Baumeister, C. and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109(5):1873–1910.

Beck, R. and Rahbari, E. (2008). Optimal reserve composition in the presence of sudden stops: The euro and the dollar as safe haven currencies. ECB Working Paper 916, European Central Bank.

Beck, R. and Rahbari, E. (2011). Optimal reserve composition in the presence of sudden stops. *Journal of International Money and Finance*, 30(6):1107–1127. DOI 10.1016/j.jimonfin.2011.02.006.

Ben-Bassat, A. (1980). The optimal composition of foreign exchange reserves. *Journal of International Economics*, 10(2):285–295.

Borio, C., Galati, G., and Heath, A. (2008). FX reserve management: Trends and challenges. BIS Papers 40, Bank for International Settlements.

Boz, E., Casas, C., Georgiadis, G., Gopinath, G., Le Mezo, H., Mehl, A., and Nguyen, T. (2022). Patterns of invoicing currency in global trade: New evidence. *Journal of International Economics*, 136:103604.

Campbell, J. Y., Serfaty-de Medeiros, K., and Viceira, L. M. (2010). Global currency hedging. *Journal of Finance*, 65(1):87–121.

Central Bank of Sri Lanka (2026). External sector statistics. Statistical releases, Central Bank of Sri Lanka.

Claessens, S. and Varangis, P. (1991). Hedging crude oil imports in developing countries. Policy Research Working Paper 755, World Bank.

Daniel, J. A. (2001). Hedging government oil price risk. IMF Working Paper 01/185, International Monetary Fund.

Erb, C. B. and Harvey, C. R. (2013). The golden dilemma. *Financial Analysts Journal*, 69(4):10–42.

European Central Bank (2023). The euro area current account after the pandemic and energy shock. ECB Economic Bulletin 6/2023, European Central Bank.

Eurostat (2026). Balance of payments statistics: Current account vintages. Statistical database, Eurostat.

Ferranti, M. (2025). Hedging sanctions risk: Cryptocurrency in central bank reserves. *Journal of International Money and Finance*, 159:103433.

Food and Agriculture Organization (2026). FAO food price index. Statistical database, Food and Agriculture Organization of the United Nations.

Gopinath, G., Boz, E., Casas, C., Díez, F. J., Gourinchas, P.-O., and Plagborg-Møller, M. (2020). Dominant currency paradigm. *American Economic Review*, 110(3):677–719.

Gürkaynak, R. S., Sack, B., and Wright, J. H. (2007). The U.S. treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, 54(8):2291–2304.

Hamilton, J. D. (2013). Historical oil shocks. In Parker, R. E. and Whaples, R., editors, *Routledge Handbook of Major Events in Economic History*, pages 239–265. Routledge.

Independent Evaluation Office of the IMF (2012). International reserves: IMF concerns and country perspectives. Ieo evaluation report, International Monetary Fund.

International Energy Agency (2021). Ammonia technology roadmap: Towards more sustainable nitrogen fertiliser production. Technical report, International Energy Agency.

International Energy Agency (2026). IEA member countries' collective emergency stock commitments. IEA statements and Oil Market Report updates, 2026.

International Monetary Fund (2015). Assessing reserve adequacy—specific proposals. Imf policy paper, International Monetary Fund.

International Monetary Fund (2023a). Pakistan: Request for a stand-by arrangement. IMF Country Report 23/260, International Monetary Fund.

International Monetary Fund (2023b). Sri lanka: Request for an extended arrangement under the extended fund facility. IMF Country Report 23/116, International Monetary Fund.

International Monetary Fund (2025). Ghana: Fourth review under the arrangement under the extended credit facility. IMF Country Report 25/175, International Monetary Fund. The Fifth Review (December 2025) reports the Gold-for-Oil programme losses and discontinuation.

International Monetary Fund (2026a). Currency composition of official foreign exchange reserves (COFER). Statistical database, International Monetary Fund.

International Monetary Fund (2026b). International financial statistics: International liquidity. Statistical database, International Monetary Fund.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review*, 111(4):1092–1125.

Khan, M. S. and Knight, M. D. (1988). Import compression and export performance in developing countries. *Review of Economics and Statistics*, 70(2):315–321.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069.

Ledoit, O. and Wolf, M. (2004). Honey, I shrunk the sample covariance matrix. *Journal of Portfolio Management*, 30(4):110–119.

Levi, P. and Molnar, G. (2022). How the energy crisis is exacerbating the food crisis. IEA Commentary, International Energy Agency.

Lu, Y. and Wang, Y. (2019). Determinants of currency composition of reserves: A portfolio theory approach with an application to RMB. IMF Working Paper 19/52, International Monetary Fund.

McCauley, R. N. and Chan, T. (2014). Currency movements drive reserve composition. *BIS Quarterly Review*, pages 23–36.

McDowell, D. (2021). Financial sanctions and political risk in the international currency system. *Review of International Political Economy*, 28(3):635–661.

Ministry of Economy, Trade and Industry, Japan (2026). Releases from national petroleum stockpiles. Press releases, March–April 2026, Ministry of Economy, Trade and Industry.

Ministry of Finance, Japan (2026). Foreign exchange intervention operations. Monthly and quarterly releases, Ministry of Finance, Japan.

Ready, R. C. (2018). Oil prices and the stock market. *Review of Finance*, 22(1):155–176.

U.S. Department of the Treasury (2026). Statements on foreign exchange operations and the FIMA repo facility. Press statements, 2026.

Vaccaro-Grange, E. (2026). Optimal currency basket estimation. IMF Working Paper 26/131, International Monetary Fund.

Working, H. (1960). Note on the correlation of first differences of averages in a random chain. *Econometrica*, 28(4):916–918.

World Bank (2020). Central bank reserve management practices: Insights into public asset management from the second RAMP survey. Technical report, World Bank.

World Bank (2023). Reserve management survey report 2023: Insights into public asset management. Technical report, World Bank.

World Bank (2026). Commodity markets: Monthly price data ("Pink Sheet"). World Bank Commodity Markets database.

Yin, L. and Han, L. (2014). Strategic allocation of commodities: A perspective based on national utility and risk hedging. *Guanli Shijie (Management World)*, (7):39–51. In Chinese; title verified against the original.

Zhang, B., Wang, X., and Hua, X. (2010a). Nominal and real returns on China's foreign exchange reserves. *Jingji Yanjiu (Economic Research Journal)*, (10):115–128. In Chinese; title verified against the original. The 2014 Social Sciences in China translation corresponds to the 2012 companion paper, not this article.

Zhang, H. (2011). *Research on the Commodity Asset Allocation of China's Foreign Exchange Reserves: An Analysis Based on Industrial Demand*. China Social Sciences Press, Beijing. In Chinese; title verified against the catalog record. ISBN 978-7-5004-9810-0; 179 pp.

Zhang, H. and Chai, Z. (2010). Foreign exchange reserve commodity asset allocation and the establishment of strategic resource reserves. *Tongji yu Juece (Statistics and Decision)*, (22). In Chinese; title and venue verified against the CNKI issue table of contents and the Wanfang record, August 2026. Page numbers not recorded in the database listings.

Zhang, H., Wu, C., and Zou, P. (2009). Physical asset allocation of foreign exchange reserve based on industrial demand. *Caijing Yanjiu (Journal of Finance and Economics)*, 35(4):4–14. In Chinese; the journal's own English title. DOI 10.3969/j.issn.1001-9952.2009.04.001.

Zhang, H., Zou, P., and Yao, W. (2010b). Foreign exchange reserve asset allocation and industrial demand. *Tongji yu Juece (Statistics and Decision)*, (22):129–132. In Chinese; title verified against the original.