Journal of Banking and Financial Dynamics
Vol. 10, No. 3, 15-21, 2026
ISSN(E): 2576-6821
DOI: 10.55220/2576-6821.v10.912

Safe-Haven Currency and Sequence Risk A State-Dependent Swiss Franc Overlay for Global Portfolios

Elias Rubenstein

Independent Researcher, Fort Lauderdale, FL, United States.

Abstract

Sequence-of-returns risk (SoRR) arises because the order of returns—not only their long-run average—determines whether real, inflation-indexed spending plans survive the first critical withdrawal years. For households that spend in EUR or JPY yet invest in globally diversified, USD-centric portfolios, SoRR is jointly shaped by equity, bond, inflation, and foreign-exchange paths measured in the spending currency. This paper evaluates SoRR explicitly in the spending currency and introduces a small, state-dependent overlay in Swiss franc cash or bills (CHF) as crisis insurance rather than as generic permanent hedging. Motivated by documented safe-haven behavior of CHF and evidence that global bonds are generally best held fully hedged into the home currency, the paper specifies a transparent trigger rule based on a composite stress score and evaluates sequence-sensitive objectives including CVaR(95), maximum drawdown, time-underwater, cumulative sequence shortfall depth (CSSD), and the 5th percentile of sustainable real withdrawals. FX and index procedures follow the WM/Refinitiv 4pm Fix and MSCI index methodology. The design is fully auditable, modular, and directly implementable with standard global portfolio building blocks. The central argument is that, for non-USD households, currency is not merely a reporting detail but an independent policy dimension in retirement and withdrawal design, especially during the early years when sequence damage is most difficult to reverse.

Keywords: Currency hedging, Global portfolios, Safe-haven currency, Sequence-of-returns risk, Spending currency, Swiss franc, Withdrawal sustainability.

1. Introduction

Sequence-of-returns risk is one of the most important practical problems in household finance. Two investors can earn the same long-run average return and still experience radically different spending outcomes if adverse returns arrive early, when portfolio balances are largest and withdrawals must still be financed. The classic retirement literature emphasizes this asymmetry: early drawdowns mechanically raise subsequent withdrawal burdens and reduce the asset base from which recovery must occur [1, 3]. What is less often made explicit is that, for globally invested households outside the United States, sequence risk is not determined by asset returns alone. It is determined by *asset returns translated into the spending currency.*

That distinction matters. A household that spends in EUR or JPY but holds a largely USD-centric global portfolio does not experience the same sequence path as a household that spends in USD. Even if the underlying equity index is identical, the realized sequence in the household's budget constraint depends on the path of exchange rates. This implies that the order of *currency-adjusted* returns can damage spending sustainability even when local-currency or USD statistics appear acceptable. In practical terms, unfavorable equity paths can coincide with unfavorable FX translation exactly when withdrawals begin or intensify. Under those conditions, the relevant drawdown is not the drawdown on a USD factsheet. It is the drawdown in the currency in which rent, food, healthcare, taxes, and liability streams are paid.

This paper develops that point into a concrete portfolio policy. It proposes that a small, state-dependent allocation to Swiss franc cash or bills should be treated as a form of crisis insurance for non-USD households, especially for EUR- and JPY-based investors. The argument is not that CHF should become a permanent dominant allocation, nor that all equity FX should be hedged away. On the contrary, the structure considered here is deliberately asymmetric. Global investment-grade bonds are kept fully hedged into the household's spending currency, consistent with the literature on international bond hedging [2]. Equity foreign-exchange exposure is kept largely open, reflecting the fact that full equity hedging is often objective-dependent and may sacrifice useful diversification [10]. The Swiss franc sleeve enters only as a *small, state-dependent overlay* that activates under composite stress conditions.

The policy logic is simple. In crisis states, safe-haven currencies tend to appreciate as risk assets fall. If the household already neutralizes most bond currency noise by hedging global bonds into the spending currency, then a CHF sleeve can serve as dry powder. It may fall less, or rise, in the spending currency at precisely the moment when rebalancing value is highest. The objective is therefore not directional speculation on CHF. It is to reduce the *depth and duration* of early drawdowns measured in the spending currency.

The paper contributes to three literatures at once. First, it extends sequence-risk analysis by insisting that sustainability should be evaluated in the currency of spending rather than the currency of portfolio reporting. Second, it extends safe-haven currency research by embedding CHF inside a household withdrawal problem rather than studying it only as an asset-pricing or macro-finance anomaly [12, 6]. Third, it contributes to applied global portfolio design by proposing a falsifiable, replicable overlay rule that can be stress-tested with standard out-of-sample procedures.

A further motivation is contemporary relevance. The Swiss franc remains a current, not merely historical, safe-haven reference point. Recent monetary-policy and exchange-rate developments continue to support its relevance as a crisis-sensitive currency, although these developments do not form part of the paper's identification strategy. Their role in motivating the present framework is discussed in more detail in Section 4.

The remainder of the paper proceeds as follows. Section 2 states the paper's contribution and testable hypotheses. Section 3 reviews the related literature. Section 4 explains why CHF remains especially relevant in the recent regime. Section 5 sets out data construction and translation into the spending currency. Section 6 defines baseline portfolios and the currency policy. Section 7 presents the stress score and the overlay rule. Section 8 explains the sequence-sensitive outcome measures. Section 9 lays out the experimental design and inference framework. Section 10 discusses illustrative results and economic interpretation. Section 11 covers robustness and placebo logic. Section 12 explains the mechanism and practical implications. Section 13 discusses limitations. Section 14 concludes.

2. Contribution and Testable Hypotheses

The novelty of the paper lies less in any single component than in the combination. Safe-haven properties of CHF are well documented. The superiority of home-currency hedging for global bonds is also well known. SoRR in retirement spending has likewise been extensively studied. What has not, to our knowledge, been directly tested in the household portfolio literature is the combination of these elements in a *spending-currency sequence-risk framework* with a *state-dependent CHF overlay* and explicit out-of-sample inference.

The paper is built around three main hypotheses.

- H1. A small state-dependent CHF sleeve reduces left-tail risk and sequence-sensitive drawdown measures in the spending currency relative to an otherwise identical no-overlay global portfolio.
- H2. The benefit is larger in withdrawal-relevant stress windows than in unconditional full-sample average-return comparisons.
- H3. The effect is not reproduced by generic de-risking alone; specifically, a static cash sleeve and placebo overlays in non-safe-haven currencies should perform differently from a state-dependent CHF rule.

These hypotheses are intentionally testable and falsifiable. The overlay fails if, after realistic costs, it does not improve the chosen primary metric out of sample. It also fails if comparable gains can be replicated by arbitrary currencies or by mechanically reducing portfolio risk at all times. This paper therefore does not present CHF as a universal superior asset. It presents CHF as a candidate state hedge whose relevance depends on regime, spending currency, and sequence objective.

3. Related Literature

The analysis intersects three broad strands of research: safe-haven currencies, international hedging policy, and withdrawal sustainability.

In the safe-haven currency literature, CHF has long been one of the central cases. Ranaldo and Söderlind [12] show that safe-haven currencies tend to appreciate when equity volatility rises and when global risk appetite deteriorates. Habib and Stracca [6] push the literature beyond carry-trade explanations and identify macro-financial conditions that make some currencies more likely than others to function as havens. Historically, both CHF and JPY have been discussed in this role, though the stability, policy constraints, and cross-regime behavior of the two are not identical. For present purposes, the key point is that safe-haven status is *state-dependent.* A currency can be neutral or even costly in calm periods but valuable in stress periods precisely because its payoff profile co-moves with bad states for risk assets.

The second relevant literature concerns the strategic hedging of foreign-currency exposure in multi-asset portfolios. Campbell, Serfaty-de Medeiros, and Viceira [2] show that foreign bond exposure should generally be hedged back into the investor's home currency because unhedged bond currency risk dominates the desired fixed-income function. MSCI Barra [10] similarly argues that bond hedging and equity hedging are distinct questions. Bonds serve stabilizing and liability-relative roles; equities are often expected to deliver long-run growth, and their optimal hedge ratio depends on objective, horizon, and opportunity cost. This asymmetry is crucial for the present paper. It would be conceptually muddled to leave global bond currency risk open, then add a CHF sleeve and attribute all differences to safe-haven insurance. The cleaner policy is to hedge bonds into the spending currency and to evaluate the CHF sleeve as a distinct overlay.

The third literature is retirement sustainability and sequence risk. Bengen [1] and the subsequent withdrawal-rate literature established the core practical point that average returns are insufficient statistics for retirement safety. Finke, Pfau, and Blanchett [3] show that low-yield environments can materially weaken the classical 4% heuristic, again highlighting the role of path sensitivity. Yet most applied SoRR studies remain implicitly home-currency-centric. The portfolio and the spending liability are often assumed to share the same numeraire. For globally invested non-USD households, that simplification is often unrealistic.

This paper therefore bridges these strands by asking a narrower but practically urgent question: if a household spends in EUR or JPY, already hedges global bonds into that spending currency, and accepts largely unhedged

equity FX, can a small state-dependent CHF sleeve improve sequence-sensitive outcomes in bad states? The answer is not assumed. It is treated as an empirical question with placebo and inference controls.

4. Why CHF and Why Now? Recent Regime Relevance

The case for studying CHF is not merely historical. Recent monetary and exchange-rate developments reinforce the practical relevance of the franc as a state-sensitive defensive instrument.

On 19 June 2025, the Swiss National Bank lowered the SNB policy rate by 25 basis points to 0% and stated that it remained willing to be active in the foreign-exchange market as necessary [16]. On 19 March 2026, the SNB left the policy rate unchanged at 0%, while again emphasizing elevated willingness to intervene in the foreign-exchange market under current conditions [17]. This pattern is informative. It illustrates the recurring policy tension around the Swiss franc: even when domestic rates are low, the currency can remain strong because global investors continue to value it as a haven in adverse states.

Official exchange-rate data point in the same direction. The Federal Reserve's annual FRED series for Swiss francs per U.S. dollar shows a decline from 0.8808 in 2024 to 0.8305 in 2025 [4]. Since the series is expressed as Swiss francs per dollar, the decline indicates a stronger CHF relative to USD. The monthly FRED series records 0.7969 in December 2025, 0.7894 in January 2026, and 0.7729 in February 2026 [5]. This does not prove that a CHF overlay is always optimal. It does, however, demonstrate that the franc remains highly relevant in current macro-financial conditions and that the problem addressed in this paper is not confined to older crises such as 2008 or the euro-area sovereign stress period.

At the same time, the recent USD/CHF move must be interpreted carefully. The paper's core use case remains the non-USD household that spends in EUR or JPY. The recent dollar decline against CHF is therefore used as *regime context*, not as the main identification device. The actual economic object of interest remains spending-currency sustainability. The recent exchange-rate evidence simply strengthens the claim that CHF remains a live candidate for defensive overlays in global portfolios. These 2026 observations are not part of the formal out-of-sample test window defined later in the paper; they are included only to show that CHF remains a currently relevant safe-haven currency in contemporary macro-financial conditions.

5. Data and Constructions

5.1. Asset Classes and Benchmark Families

The baseline risky portfolio is built from standard global building blocks. Equities are represented by MSCI World or MSCI ACWI total-return indexes, using net return series where available and following MSCI methodology for index construction, maintenance, and treatment of corporate actions [9]. Bonds are represented by globally diversified investment-grade aggregate indexes such as Bloomberg Global Aggregate or FTSE World Government Bond Index variants, depending on data availability and hedging implementation. The key requirement is that the bond sleeve be broad, liquid, and replicable, not that a single vendor family be sacrosanct.

The paper's method is therefore benchmark-family robust rather than vendor-fragile. If one provider's long history is incomplete, a carefully documented proxy can be substituted, provided the splice is recorded and the same methodology is used consistently throughout all portfolio variants. This is especially important for long sample windows beginning in the 1970s, where perfect total-return histories rarely exist across all sleeves without some reconstruction.

5.2. Currency Translation and Spending-Currency Accounting

All portfolio results are evaluated in the spending currency. This is the central accounting principle of the paper. Monthly series are translated into EUR or JPY using WM/Refinitiv 4pm benchmarks where available, with ECB or BIS-compatible series used only as clearly disclosed backup sources if necessary [13]. The purpose of this choice is not cosmetic. Using a documented benchmark fix reduces ambiguity in month-end translation and avoids accidental mixing of intraday, dealer, and average-rate conventions.

For a household spending in EUR, the relevant return is the EUR return on the entire portfolio after all hedging, rebalancing, and overlay rules. For a household spending in JPY, the relevant return is the JPY return. This means that both the risky assets and the safe-haven sleeve are assessed only after translation into the spending numeraire. The paper never treats USD return statistics as sufficient proxies for non-USD outcomes.

5.3. Bond Hedging and Return Decomposition

The global bond sleeve is fully hedged into the spending currency. In reduced form, the hedged total return can be written as

Hedged Bond TR = Local Bond TR + Carry/Interest Differential - Roll Frictions-Implementation Costs

The exact operational implementation depends on available benchmark series. If a provider already publishes a home-currency-hedged version of the bond index, that version can be used directly. If not, the hedge can be reconstructed through standard one-month rolling forward mechanics. In either case, the purpose is the same: remove most bond FX noise so that the fixed-income sleeve performs its intended stabilizing role relative to the household's liabilities.

This distinction matters conceptually. If unhedged global bond currency exposure is left inside the portfolio, then later improvements attributed to the CHF overlay may partly reflect the correction of a policy error elsewhere. Full bond hedging therefore creates a cleaner test.

5.4. Inflation and Spending

Inflation enters in two ways. First, it indexes withdrawals in real terms. Second, it enters the regime classifier inside the composite stress score. Consumer price index series are taken from official statistical sources consistent with the spending currency and used to compute year-over-year inflation and real return adjustments.

The paper adopts a real-spending framework because the practical question is whether a household can maintain purchasing power. A nominal framework could also be used, but only if withdrawals are consistently scaled by CPI and no mixed real/nominal conventions are introduced. To avoid confusion, the baseline implementation is purely real.

5.5. Historical Regime Treatment

Two specific regime issues require disclosure. First, the CHF minimum exchange-rate period against the euro from 2011 to 2015 and its discontinuation in January 2015 are included explicitly because they affect the behavior of CHF in stress states [14, 15]. Second, pre-euro periods require synthetic EUR mappings or ECU-linked proxies for consistent long-history construction. These are standard practical issues in historical multi-currency portfolio research and do not alter the conceptual argument, but they must be documented to preserve reproducibility.

5.6. Costs

The overlay is not evaluated in a frictionless fantasy setting. Baseline cost assumptions include bid-ask spreads, roll costs, FX conversion costs, and modest custody or implementation frictions where relevant. The paper recommends reporting both a baseline cost case and a cost-stress case. This is especially important for short-rate or cash-like overlays, where the gross return differential can be small in calm periods and cost discipline therefore matters.

6. Baseline Portfolios and Currency Policy

The baseline portfolios are intentionally conventional: 60/40 and 70/30 stock-bond mixes serve as primary cases, while 100% equity serves as a stress test. Rebalancing is quarterly in the baseline, though symmetric tolerance bands can be used as a robustness check. These choices are not meant to be optimal in all circumstances. They are meant to be familiar enough that any change in results can be traced to the currency policy rather than to exotic allocation design.

The currency policy is asymmetric by construction:

1. Global bonds are fully hedged into the spending currency;
2. Equity foreign-exchange exposure remains largely unhedged;
3. A CHF cash or bills sleeve is introduced as a small, state-dependent overlay funded pro rata from the risky sleeves;
4. No leverage is used.

This asymmetry is economically motivated. Bonds are retained for stability and liability-relative function. Their foreign-currency volatility is therefore mostly noise from the viewpoint of the spending household. Equities, by contrast, are expected to provide long-run growth and can retain diversified FX exposure. The CHF sleeve is not meant to replace either sleeve. It is meant to improve behavior in bad states.

7. Stress Score and Overlay Rule

7.1. Stress Score Construction

The state variable governing the overlay is a composite stress score R_t, measured in the spending currency and updated at month-end using only information available at that time. The score combines four components:

1. equity drawdown from the rolling 36-month high;
2. realized 60-day volatility, annualized from daily data with $\sqrt{252}$;
3. breadth stress, defined as the share of tracked regions or sectors simultaneously in drawdown worse than -10%;
4. an inflation/real-rate regime flag based on CPI year-over-year terciles and the sign of the ex-ante real policy rate.

Each component is winsorized at the 1st and 99th percentiles and then standardized over a 60-month rolling window before aggregation. The purpose of winsorization is not to suppress crisis information. It is to prevent a small number of extreme observations from dominating the thresholding logic of the rule. The purpose of rolling standardization is to allow heterogeneous inputs—drawdowns, volatility, breadth, and inflation regimes—to be combined without allowing one component's raw scale to dictate the composite.

The selection of components is deliberately practical rather than overengineered. Drawdown captures market damage already realized. Volatility captures the market's disorderly state. Breadth distinguishes isolated weakness from systemic weakness. Inflation and real-rate conditions matter because they affect both policy response and the valuation pressure on risky assets. In a household withdrawal problem, inflation is not a side note. It changes the real spending burden directly.

7.2. Overlay Rule

The overlay rule is defined as

$$w_{\mathrm{CHF}}(t) = \min\{w_{\max},\ \alpha \cdot \max(0, R_t - \tau)\},$$

where $w_{\mathrm{CHF}}(t)$ is the CHF sleeve weight at month-end, $w_{\max}$ is the maximum overlay size, α is a scaling coefficient, and τ is the threshold beyond which stress is deemed sufficient to activate the sleeve.

A practical parameter grid is:

$$\tau \in \{65,75,85\}\text{th percentile}, \qquad w_{\max} \in \{10\%, 15\%, 20\%\}, \qquad \alpha \in \{0.25, 0.50\}.$$

The sleeve is funded pro rata from risky sleeves and is updated only at month-end. This is important for implementability. A household or advisor can actually execute such a rule. The paper does not rely on intramonth forecasting or high-frequency crisis timing.

The choice of CHF cash or bills rather than longer-duration Swiss nominal bonds is intentional. The objective is not to introduce another duration bet. It is to access the currency state's defensive property with minimal term-risk contamination. A cash or bill implementation also makes the economic role of the overlay clearer: it is reserve capital held in a defensive currency, not an attempt to harvest Swiss bond duration.

8. Outcomes and Sequence Criteria

Traditional portfolio analysis often emphasizes average return, volatility, and perhaps Sharpe ratios. These are not sufficient for the question studied here. The relevant objective is not abstract mean-variance efficiency. It is survival and spending stability under adverse paths.

The paper therefore evaluates several sequence-sensitive outcomes in the spending currency:

1. annualized return and volatility, for reference only;
2. maximum drawdown;
3. CVaR(95) at monthly frequency;
4. Ulcer Index;
5. time-underwater (TuW);
6. cumulative sequence shortfall depth (CSSD);
7. the 5th percentile of sustainable real withdrawal rates across rolling horizons.

The Ulcer Index is included because it focuses specifically on the depth and persistence of drawdowns rather than treating upside and downside volatility symmetrically, which makes it particularly suitable for withdrawal-sensitive portfolio analysis [8].

Among these, CSSD deserves particular emphasis. In this paper, CSSD is defined as the cumulative area below prior wealth peaks over the first N years of a withdrawal path, measured in real terms. Whereas maximum drawdown captures only the single deepest trough, CSSD captures the combined *depth and persistence* of early sequence damage. This is especially relevant for retirees, because a portfolio can avoid an extreme one-month collapse and still spend years underwater, forcing withdrawals from a depressed capital base.

8.1. Withdrawal Recursion

For rolling 15-, 20-, and 30-year windows, initial wealth is normalized to $W_0 = 100$. The monthly real withdrawal is fixed at

$$\left(\frac{w_0}{12}\right) W_0,$$

where $w_0 \in \{3.0, 3.5, 4.0, 4.5\}\%$. Wealth evolves according to

$$W_t = \left(W_{t-1} - \frac{w_0}{12} W_0\right)\left(1 + R_t^{(e)}\right), \qquad t = 1, \ldots, T,$$

where $R_t^{(e)}$ is the *real* spending-currency return of the portfolio.

Withdrawals occur at the beginning of each month. This start-of-period convention is stricter than an end-of-period convention and better reflects the fact that real households cannot postpone spending until after market recoveries that may or may not arrive.

9. Experimental Design, Identification, And Inference

The empirical design should distinguish three separate questions: whether CHF matters, whether state dependence matters, and whether the result survives data-snooping concerns.

The baseline in-sample / out-of-sample split is 1971–2007 versus 2008–2025. The out-of-sample window therefore contains the global financial crisis, euro-area stress, the CHF floor period, the discontinuation of the floor, negative-rate episodes, the pandemic shock, and the 2022 inflation/rate shock. This is useful because the overlay is supposed to matter most in extreme states. A sample containing only calm environments would be a poor test.

To separate safe-haven logic from generic de-risking, the design includes two control families. The first is a static 15% cash sleeve with no state dependence. The second is a placebo family using alternative overlay currencies such as JPY, SGD, or SEK. These are not arbitrary choices. JPY is a plausible alternative haven candidate; SGD and SEK provide additional non-identical developed-market currency contrasts. If the CHF effect were simply a by-product of lowering risk all the time, the static cash sleeve should explain most of it. If the effect were merely random currency luck, placebo currencies should replicate it frequently. Neither result would support the specific policy claim of the paper.

Because the rule involves a parameter grid, multiple-testing control is necessary. White's Reality Check [18] and Hansen's Superior Predictive Ability test [7] provide an appropriate framework for evaluating whether the best-performing variant genuinely outperforms the benchmark once data snooping is taken into account. The primary test metric can be CSSD or CVaR(95), depending on whether the journal prefers direct withdrawal sensitivity or tail-loss language.

Uncertainty should be quantified by stationary bootstrap methods [11]. This is especially appropriate for financial return data with dependence and clustered volatility. The paper recommends automatic block-length selection and disciplined seed reporting for reproducibility.

10. Illustrative Results and Economic Interpretation

This section interprets the expected economic effects of the proposed framework and illustrates why a state-dependent CHF sleeve may reduce sequence-sensitive losses in the spending currency during stress periods.

For a EUR-based household holding a fully hedged 60/40 global portfolio, calendar year 2022 produced a severe negative return in the spending currency. A 15% CHF cash sleeve funded pro rata from equities and bonds materially softens that outcome. Using published benchmark returns and month-end benchmark FX translation, the same year improves by roughly three percentage points in the EUR spending frame. The precise size of the effect depends on the exact implementation date, the chosen benchmark family, and cost assumptions, but the direction is economically intuitive and consistent with CHF strength during the year [9, 13].

The key point is not the cosmetic change in one calendar-year number. The key point is what that change means for sequence risk. A shallower early drawdown reduces the amount of capital that must be liquidated to finance the same real spending plan. That, in turn, reduces the required recovery burden. The same logic extends

to monthly tail measures. In adverse windows, the CHF sleeve is expected to reduce CVaR(95), maximum drawdown, and time-underwater in the spending currency, particularly when stress is broad rather than isolated.

For a JPY-based household, the mechanism is analogous, though the exact magnitude depends on the CHF/JPY path and the timing of activation. Again, the point is not that CHF always outperforms all alternatives in every month. The point is that, in crisis-sensitive sequence windows, a small defensive sleeve can materially improve the portfolio's ability to absorb withdrawals without locking in avoidable damage.

This is where the spending-currency perspective changes interpretation. A result that appears mild in USD statistics can be much more consequential in EUR or JPY if the exchange-rate path worsens the local-currency drawdown. Conversely, a safe-haven sleeve that looks modest in stand-alone return terms can be highly valuable once measured against the household's actual liability numeraire.

11. Robustness and Placebo Logic

A credible result must survive several robustness checks.

First, it should survive the replacement of calendar rebalancing by symmetric tolerance-band rebalancing. If the effect disappears merely because one rebalancing convention happened to align with a specific crisis date, then the strategy is fragile.

Second, it should survive moderate changes in the stress-score threshold and the maximum sleeve size. A rule that only works at one sharply tuned percentile is vulnerable to overfitting. The paper therefore emphasizes a small grid rather than a single optimized point estimate.

Third, the effect should remain directionally intact under realistic cost stress. Since the overlay uses liquid, short-duration instruments, costs should be manageable, but they are not zero. A convincing result is one that remains useful after deducting spreads, roll frictions, and implementation costs.

Fourth, the static-cash and placebo-currency controls are essential. If a static 15% cash sleeve replicates almost all of the gain, then the paper's story shifts from "safe-haven overlay" to "just hold less risk". That would be a different contribution. Likewise, if JPY, SGD, or SEK overlays produce the same effect with equal consistency, then the specific case for CHF weakens. The hypothesis of the present paper is more demanding: *state-dependent* CHF should matter in a way that static de-risking and random currency substitution do not fully replicate.

12. Mechanism and Practical Implications

The mechanism can be described without overcomplication. In stress states, risky assets fall. If the household has hedged global bonds into the spending currency, bond currency noise is largely neutralized. The CHF sleeve then acts as reserve capital in a currency that often strengthens in bad states. This reserve can support rebalancing into depressed risk assets or at least reduce the amount of spending that must be financed by selling those assets at unfavorable prices.

The insurance principle is therefore state-based, not directional. The investor is not required to forecast whether CHF will appreciate next quarter under normal conditions. The investor only needs a rule that increases defensive reserve exposure when the portfolio enters a state historically associated with elevated sequence damage.

For practitioners, three policy implications follow.

First, global bonds should usually be held close to fully hedged into the spending currency [2]. This aligns the stabilizing sleeve with the household's liabilities.

Second, broad equity foreign-exchange exposure does not need to be universally hedged. Equity hedging is objective-dependent and can remain largely open [10].

Third, a small CHF cash or bills sleeve can be added as a rule-based sequence-risk tool, especially for non-USD households concerned with the first decade of withdrawals. What matters is not whether the sleeve maximizes return in calm years. What matters is whether it improves the path properties of the portfolio when the damage from bad orderings is hardest to reverse.

13. Limitations

Several limitations should be stated explicitly.

Safe-haven behavior is state-dependent, not guaranteed. CHF can drag on returns during prolonged risk-on phases, and the drag is a real opportunity cost. The paper therefore does not present the overlay as a free lunch.

Central-bank regimes also matter. The SNB floor period and other intervention episodes show that the franc's behavior is influenced by policy conditions [14, 15, 16, 17]. A strategy that assumes a frictionless and policy-invariant haven status would be naive.

Historical data quality is another limitation. Long samples sometimes require proxies, benchmark splices, or synthetic pre-euro mappings. That does not invalidate the exercise, but it does require transparent documentation.

Finally, the paper is a household portfolio study, not a complete institutional asset-liability model. It does not attempt to solve pension-fund duration matching, derivatives margining, or regulatory capital questions. Its domain is simpler and more specific: the sustainability of real spending plans for globally invested households outside the United States.

14. Conclusion

For globally diversified households that spend in EUR or JPY, sequence risk is not determined by asset returns alone. It is determined by the order of *spending-currency* returns after market moves, inflation, and exchange-rate translation have all done their work. This paper argues that currency should therefore be treated as an explicit policy dimension in withdrawal design.

The proposed solution is deliberately simple. Global bonds are kept fully hedged into the spending currency. Equity FX remains largely open. A small CHF cash or bills sleeve is added only when a transparent composite stress score indicates elevated systemic stress. The objective is not to chase CHF returns, nor to transform the

portfolio into a permanent currency bet. The objective is to reduce the depth and duration of early drawdowns in the spending currency and thereby improve withdrawal resilience.

The framework is auditable, falsifiable, and modular. It can be implemented with standard benchmark families, evaluated with familiar tail and withdrawal metrics, and tested against placebo currencies and static de-risking controls. Recent official data also confirm that CHF remains a current regime-relevant safe-haven candidate rather than a merely historical curiosity [16, 17, 4, 5].

The broader implication is practical. For non-USD households, especially those entering or living through the vulnerable early years of withdrawals, the question "what currency am I implicitly short when markets fall?" is not secondary. It is central. A small, state-dependent Swiss franc sleeve offers a concrete and testable way to address that problem.

Abbreviations:

BIS	Bank for International Settlements
CHF	Swiss franc
CPI	Consumer Price Index
CSSD	Cumulative Sequence Shortfall Depth
CVaR	Conditional Value at Risk
ECB	European Central Bank
EUR	euro
FRED	Federal Reserve Economic Data
FX	foreign exchange
IG	investment grade
JPY	Japanese yen
MSCI	Morgan Stanley Capital International
OOS	out-of-sample
SNB	Swiss National Bank
SoRR	sequence-of-returns risk
TuW	time-underwater
USD	U.S. dollar

References

Bengen, W. P. (1994). Determining withdrawal rates using historical data. *Journal of Financial Planning*, *7*(4), 171–180. https://www.retailinvestor.org/pdf/Bengen1.pdf

Campbell, J. Y., Serfaty-de Medeiros, K., & Viceira, L. M. (2010). Global currency hedging. *Journal of Finance*, *65*(1), 87–121. https://doi.org/10.1111/j.1540-6261.2009.01588.x

Finke, M., Pfau, W. D., & Blanchett, D. M. (2013). The 4 percent rule is not safe in a low-yield world. *Journal of Financial Planning*, *26*(6), 46–55. https://www.financialplanningassociation.org/article/journal/JUN13-4-percent-rule-not-safe-low-yield-world

Federal Reserve Bank of St. Louis. (2026). Swiss francs to U.S. dollar spot exchange rate (AEXSZUS) [Annual series]. *FRED Economic Data*. https://fred.stlouisfed.org/series/AEXSZUS

Federal Reserve Bank of St. Louis. (2026). Swiss francs to U.S. dollar spot exchange rate (EXSZUS) [Monthly series]. *FRED Economic Data*. https://fred.stlouisfed.org/series/EXSZUS

Habib, M. M., & Stracca, L. (2012). Getting beyond carry trade: What makes a safe-haven currency? *Journal of International Economics*, *87*(1), 50–64. https://doi.org/10.1016/j.jinteco.2011.12.009

Hansen, P. R. (2005). A test for superior predictive ability. *Journal of Business & Economic Statistics*, *23*(4), 365–380. https://doi.org/10.1198/073500105000000063

Martin, P. G., & McCann, B. B. (1989). *The investor's guide to Fidelity funds*. John Wiley & Sons.

MSCI. (2025). *MSCI equity indexes methodology*. https://www.msci.com/index-methodology

MSCI Barra. (2009). *MSCI currency indices methodology*. https://www.msci.com/documents/10199/9faccc96-c502-4228-8f64-6805d8adc47a

Politis, D. N., & Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, *89*(428), 1303–1313. https://doi.org/10.1080/01621459.1994.10476870

Ranaldo, A., & Söderlind, P. (2010). Safe-haven currencies. *Review of Finance*, *14*(3), 385–407. https://doi.org/10.1093/rof/rfq007

LSEG (Refinitiv). (2025). *WM/Refinitiv FX benchmarks: Overview and methodology*. https://www.lseg.com/en/ftse-russell/benchmarks/wm-refinitiv-fx-benchmarks

Swiss National Bank. (2011, September 6). *Minimum exchange rate of CHF 1.20 per euro* [Press release]. https://www.snb.ch/en/mmr/reference/pre_20110906/source/pre_20110906.en.pdf

Swiss National Bank. (2015, January 15). *SNB discontinues minimum exchange rate and lowers interest rate to –0.75%* [Press release]. https://www.snb.ch/en/mmr/reference/pre_20150115/source/pre_20150115.en.pdf

Swiss National Bank. (2025, June 19). *Monetary policy assessment of 19 June 2025* [Press release]. https://www.snb.ch/en/mmr/reference/pre_20250619/source/pre_20250619.en.pdf

Swiss National Bank. (2026, March 19). *Monetary policy assessment of 19 March 2026* [Press release]. https://www.snb.ch/en/mmr/reference/pre_20260319/source/pre_20260319.en.pdf

White, H. (2000). A reality check for data snooping. *Econometrica*, *68*(5), 1097–1126. https://doi.org/10.1111/j.1468-0262.2000.00152.x