2026

Research Award

icpm | INTERNATIONAL CENTRE FOR PENSION MANAGEMENT

icpm | INTERNATIONAL CENTRE FOR PENSION MANAGEMENT

Contents

1. Dollar Asset Holdings and Hedging Around the Globe

Authors:

Amy Huber, The Wharton School of the University of Pennsylvania
Wenxin Du, Harvard Business School

Key Takeaways

- Foreign investors' holdings of U.S. dollar (USD) securities grew sixfold over two decades, from about $5.5 trillion in 2002 to roughly $33.4 trillion in 2021, with foreign mutual funds, pension funds, and insurance companies materially increasing the USD share of their overall portfolios after the 2008-09 global financial crisis.
- A meaningful share of this USD exposure is hedged: by 2020, USD hedge ratios reached about 44% for insurance companies, 35% for pension funds, and 21% for mutual funds. Together these three sectors generated nearly $2 trillion of outstanding USD hedges, with hedge ratios up roughly 15 percentage points relative to pre-crisis levels.
- Hedging decisions are driven not only by variance minimization but also by expected currency returns arising from deviations from uncovered interest parity (UIP) and covered interest parity (CIP). Cross-sectional and time-series patterns in observed FX exposures are far better explained when these expected-return considerations are added to the standard mean-variance framework.
- Aggregate hedging demand materially affects hedging costs because the dealers (typically large global banks) that supply hedges face balance-sheet constraints. The cross-sectional R-squared between hedging volume and CIP deviations is 0.73, and granular data on Japanese insurers establish a causal link from hedging demand to the cross-currency basis.

The full paper is available here.

Summary

Motivation: The U.S. dollar is the dominant currency in cross-border securities holdings. But owning a USD-denominated bond or stock is not the same as bearing USD currency risk: investors can use FX forwards and swaps to hedge most of the exchange-rate exposure while still holding the underlying asset. As a result, headline figures on USD asset holdings can substantially overstate foreign investors' true dollar exposure. Despite the size of the underlying flows (FX-derivative average daily turnover reached $5.4 trillion in 2022, with 88% involving USD pairs), there is little systematic empirical evidence on how foreign institutional investors manage USD currency risk. Existing models typically assume foreign asset demand

is either fully unhedged or, for bonds, fully hedged. This paper distinguishes the demand for USD-denominated assets from the demand for USD currency exposure, providing the first comprehensive cross-country picture of foreign holdings and hedging practices.

Methodology: The authors hand-collect and harmonize a wide range of sector statistics and company filings from global institutional investors to build sector-level estimates of USD security holdings and FX hedging across seven sectors: official, banks, insurance, pension funds, mutual funds, non-financial corporations and households, and hedge funds. They focus particularly on mutual funds, pensions, and insurers — the sectors with the most actively managed currency exposure. They complement the demand-side data with hedging-supply estimates derived from BIS banking statistics and granular firm-level filings of Japanese insurers. They develop a mean-variance portfolio choice framework that augments the standard variance-covariance motive with expected FX returns from UIP and CIP deviations, derive predictions for how hedge ratios should respond, and test these predictions in the data. To examine the link between aggregate hedging and FX-derivative pricing, they exploit cross-currency variation in CIP deviations and use a granular instrumental-variables strategy on Japanese insurer data to causally identify the effect of hedging demand on the cross-currency basis.

Findings:

- Foreign USD security holdings grew from $5.5 trillion in 2002 to about $33.4 trillion in 2021. The increase reflects more than wealth growth: post-GFC, mutual funds, insurers, and pensions raised the USD share of their overall portfolios by roughly 7.7 percentage points and the USD share of their non-domestic investments by 6.6 percentage points.
- Higher USD asset holdings did not translate one-for-one into higher USD currency exposure. Hedge ratios in actively managed sectors rose substantially after the GFC, by an average of about 14.7 percentage points. By 2020, USD hedge ratios were about 44% (insurance), 35% (pensions), and 21% (mutual funds).
- Elevated and volatile CIP deviations made hedging costlier post-GFC. The authors estimate that, between 2017 and 2020, hedging costs from short-term CIP deviations averaged about $2.7 billion per annum for the insurance and pension sectors alone.
- Hedging practices vary materially across investors, currency areas, and banking systems. There is suggestive evidence that bonds are hedged at higher ratios than equities, but USD hedge ratios still vary considerably even within the same sector.
- The authors' mean-variance model, which augments variance-covariance with expected FX returns from UIP and CIP deviations, reconciles observed FX exposures markedly better than the variance-only benchmark in both the cross-section and the time series. Sectoral differences in hedging align with differences in effective risk tolerance.
- On the supply side, global banks supply substantial volumes of FX hedges through their balance sheets, with significant cross-country heterogeneity: some banking systems are net suppliers of USD FX hedges, while others use FX hedging primarily to fund their own USD operations.
- Using granular Japanese-insurer data and a granular instrumental-variables strategy, the authors find that a $1 billion increase in USD/JPY hedging supplied by intermediaries

widens the 3-month cross-currency basis by about 0.3 bps, and a €1 billion increase in EUR/JPY hedging widens the basis by about 0.5 bps — providing causal evidence that aggregate hedging demand affects hedging costs in the presence of constrained intermediaries.

2. Efficient Decumulation Strategy with Long-term Care Insurance and Guaranteed Minimum Death Benefit

Authors:

Jennifer Alonso-Garcia, Department of Mathematics, Faculté des Sciences, Université Libre de Bruxelles
Jonathan Ziveyi, School of Risk and Actuarial Studies, UNSW Sydney
Mengdie Hu, School of Risk and Actuarial Studies, UNSW Sydney
Yuxin Zhou, School of Risk and Actuarial Studies, UNSW Sydney

Key Takeaways

- The paper proposes an integrated retirement decumulation framework that combines a state-contingent ("efficient") drawdown rule, a target-volatility investment strategy, long-term care insurance (LTCI), and a guaranteed minimum death benefit (GMDB) — addressing market, longevity, late-life health, and bequest risks within a single model rather than treating them in isolation.
- Numerical simulations on Australian retiree households show that the integrated strategy produces a smoother, gradually declining consumption path and is more resilient to adverse market shocks than rule-of-thumb minimum-withdrawal strategies and full annuitization.
- Sensitivity analyses suggest moderate target-volatility levels strike the best risk/sustainability trade-off, higher LTCI allocations meaningfully reduce unpaid late-life health costs, and inflation in aged-care costs accelerates wealth depletion — underscoring the need to calibrate investment risk, insurance coverage, and inflation assumptions jointly.
- When combined with the Australian Age Pension and full LTCI coverage, the proposed efficient drawdown rule consistently delivers higher early-retirement consumption than age-based minimum-withdrawal rules and full annuitization.

The full paper is available here.

Summary

Motivation: With the global shift from defined benefit to defined contribution pensions, longevity, market, and late-life health risks now sit on individuals' balance sheets. This has prompted excessively conservative drawdown behaviour by retirees, who tend to underspend out of precaution. Existing retirement-decumulation research falls into three strands: heuristic

withdrawal rules (simple but ignoring market and health shocks), formal optimisation (theoretically coherent but hard to calibrate and implement), and insurance/annuitisation products (each addressing one risk in isolation). What is missing is a framework that jointly handles withdrawal smoothing, portfolio risk control, late-life-care protection, and bequest objectives, while remaining transparent and practically interpretable. The authors aim to fill this gap with a unified decumulation framework grounded in interacting health, market, and insurance dynamics.

Methodology: At retirement, the household is assumed to allocate initial wealth across three components: an LTCI premium covering disability-state health expenditure, a GMDB rider securing a minimum bequest, and liquid investable wealth. The remaining liquid wealth is decumulated through a time-varying, state-contingent annuity-style withdrawal rule whose annuity factors depend on health state and remaining wealth. The investment portfolio is governed by a target-volatility strategy (TVS) that adjusts asset weights to dampen wealth-path volatility. Health dynamics follow a multi-state model with healthy, mildly disabled, fully disabled, and deceased states; the financial market follows a regime-switching model. Through Monte Carlo simulation (10,000 paths) over a multi-decade retirement horizon, the authors evaluate consumption stability, wealth sustainability, late-life-care protection, and bequest provision under interacting risks. The framework is calibrated to Australian retiree data (e.g., AUD 600,000 initial wealth, ASFA retirement-income benchmarks) and benchmarked against age-based minimum-withdrawal rules and full annuitisation, with sensitivity analyses over insurance allocation ratios, target-volatility levels, health-transition assumptions, and aged-care inflation.

Findings:

- In the baseline scenario combining LTCI, GMDB, and the target-volatility strategy, the framework produces a smoother and gradually declining consumption path. The TVS dampens transmission of market shocks into consumption, while LTCI and GMDB absorb late-life-care and bequest risks.
- Higher LTCI allocations substantially reduce the share of retirement paths in which long-term care costs go unpaid, while higher target-volatility levels increase the probability that wealth becomes insufficient to cover care needs — pointing to a clear trade-off between investment ambition and care-cost protection.
- Inflation in aged-care costs accelerates wealth depletion and concentrates large unpaid-care outcomes among paths with otherwise strong investment performance. This indicates that medical-cost inflation is a first-order risk for late-life affordability and must be modelled jointly with portfolio choice and insurance coverage.
- Cross-sectional wealth distributions exhibit a persistent upper tail driven by favourable investment and health trajectories, even as the centre of the distribution shifts down with withdrawals. Substantial heterogeneity in retirement outcomes underscores the importance of modelling financial and health risks jointly rather than separately.
- Moderate target-volatility levels deliver the most attractive balance between consumption sustainability and risk; the strategy remains robust across alternative health-transition scenarios.

- When the GMDB is removed and full-coverage LTCI is paired with the Australian Age Pension, the efficient drawdown rule consistently delivers higher consumption than both age-based minimum-withdrawal rules and full annuitisation in the early years of retirement, when retirees value spending most.
- The authors flag several extensions: jointly modelling general and medical inflation as stochastic processes, endogenising LTCI and GMDB allocation decisions (e.g., via reinforcement learning), incorporating waiting periods and reimbursement structures in LTCI contracts, and embedding a deferred annuity to provide a late-life income floor.

3. The Pricing of Geopolitical Tensions over a Century

Authors:

Andrei Goncalves, The Ohio State University
Alessandro Melone, The Ohio State University
Andrea Ricciardi, The Ohio State University

Key Takeaways

- Decomposing geopolitical risk into geopolitical threats (GPT — anticipated tensions) and geopolitical acts (GPA — realised events) using nearly a century of data reveals that the two have very different financial-market consequences. GPT, not GPA, is what investors and firms actually respond to.
- GPT closely tracks subjective measures of geopolitical risk perceptions (ICRG ratings, BofA fund-manager surveys), drives capital allocation decisions in 13F portfolio holdings, raises perceived investment risk, and predicts lower future aggregate, industry-, and firm-level investment. GPA shows weaker and less stable links to all of these.
- GPT is priced in the cross-section of U.S. equities, in cross-sections of equity anomaly portfolios, and in cross-sections of country-level equity and bond indices. The GPT risk premium remains after controlling for the CAPM, ICAPM, characteristics-based factor models, and other news-based risk indices including aggregate geopolitical risk, war discourse, economic-policy uncertainty, expected market volatility, and trade-policy uncertainty.
- GPT (but not GPA) significantly forecasts time variation in country-level equity premia, especially at three- to five-year horizons. The mechanism most consistent with the data is time-varying disaster risk; non-linear market risks and overreaction to threats receive less empirical support.

The full paper is available here.

Summary

Motivation: Geopolitical tensions — exemplified today by the Russia-Ukraine conflict and the Middle East crisis — are now viewed by central banks and sovereign wealth funds as the primary risk to global growth. A large literature documents the macroeconomic effects of geopolitical risk on production, investment, employment, consumption, inflation, innovation, trade, and oil prices. Much less is known about how geopolitical risk shapes the beliefs and capital allocation decisions of investors and firms — and through them, risk premia. Two challenges complicate this question. First, geopolitical risks reflect both realised events and

the threat of future events, which may have very different effects on forward-looking markets. Second, geopolitical tensions are infrequent and clustered in time; calm periods reduce risk premia and require a long sample with enough episodes of elevated tensions for proper inference. The authors leverage nearly a century of news-based risk indices to address both challenges.

Methodology: The authors use the news-based geopolitical risk indices of Caldara and Iacoviello (2022), constructed from the share of newspaper articles in major U.S. newspapers that discuss geopolitical conflicts. The Caldara-Iacoviello geopolitical risk (GPR) index is decomposed into a geopolitical threats (GPT) index — articles on military buildups, terrorist threats, and similar — and a geopolitical acts (GPA) index — articles on the beginning of wars, terrorist attacks, and similar. The sample spans 1930-2024. The authors first relate GPT and GPA to subjective measures of geopolitical risk (ICRG ratings, BofA fund-manager surveys), to investor capital allocation in 13F holdings, and to firm-level capital expenditure. They then estimate cross-sectional risk premia using portfolio sorts on rolling-window univariate betas of stock returns onto each risk index. For equity anomaly portfolios, they apply Supervised Principal Component Analysis (Giglio, Xiu, and Zhang, 2025). For country-level indices, they use the Jordà et al. (2019) dataset of 16 developed countries and Fama-MacBeth regressions. They study time-series predictability of country-level equity premia at horizons up to five years and compare GPT against GPR, GPA, and other news-based risk indices. Finally, they evaluate three potential mechanisms — non-linear market risks, time-varying disaster risk, and overreaction to threats — to see which is most consistent with the data.

Findings:

- GPT rises in anticipation of major geopolitical conflicts, while GPA rises once events materialise — confirming the forward-looking interpretation of GPT as geopolitical uncertainty and GPA as the realisation of geopolitical events.
- Time variation in GPT (but not GPA) closely tracks subjective geopolitical-risk measures from ICRG ratings and BofA fund-manager surveys. In 13F holdings, investors allocate less to stocks with higher GPT exposure, and the effect strengthens during periods of high attention to geopolitical risks.
- GPT increases also predict lower future aggregate and industry-level investment, while at the firm level firms with higher GPT betas systematically cut capital expenditures. GPA does not produce the same effects, and the GPT effect strengthens when attention to geopolitical risk is high.
- In the cross-section of U.S. stocks (1930-2024), portfolios sorted on GPT betas produce an economically and statistically significant risk premium that survives after controlling for the CAPM, ICAPM, all the other risk indices considered, and a broad set of characteristics-based factor models. The GPT premium is present both across and within industries, is positive throughout the sample period, and is elevated when attention to geopolitical risk is high.
- For equity anomaly portfolios (1963 onward), the GPT mimicking factor delivers a significant risk premium that survives controlling for CAPM, ICAPM, and the mimicking factors of other risk indices. The GPA and GPR mimicking factors have negative and statistically insignificant alphas after controlling for the GPT mimicking factor.

- For country-level equity and bond indices, the GPT mimicking factor again provides a significant risk premium that survives controlling for the World CAPM. In contrast, whether other risk indices (including GPA) produce significant premia varies by index, period, and asset cross-section.
- GPT (but not GPA) significantly forecasts country-level equity premia, with the strongest effects at three- and five-year horizons. The link to bond risk premia is weaker, possibly reflecting flight-to-safety and pass-through of wartime fiscal costs to bondholders via inflation and financial repression.
- Among three candidate channels, time-varying disaster risk is the mechanism most consistent with the documented GPT premium. The authors emphasise that distinguishing shocks to the level of macroeconomic conditions from shocks to uncertainty about those conditions — as the GPA/GPT decomposition naturally does — is likely useful in other contexts including monetary, fiscal, and regulatory uncertainty.

4. The Private Capital Alpha

Authors:

Andrei Goncalves, The Ohio State University
Gregory Brown, University of North Carolina at Chapel Hill
Wendy Hu, MSCI

Key Takeaways

- The traditional way of estimating private-capital alpha (regressing private-fund index returns on public-market indices) ignores key economic realities of private investing — illiquidity, NAV staleness, underdiversification, and the inability of LPs to control their effective allocation — and produces materially overstated alphas.
- Once these realities are accounted for via realistic LP simulations, the alpha of buyout (BO) is positive and statistically significant at 2.5% per annum (95% confidence interval 0.1%-4.5%) over 1987-2022, but the alpha of venture capital (VC) is large but statistically unreliable (4.5%, CI -4.4% to 15.5%), and the alpha of real estate (RE) is small and not statistically distinguishable from zero (2.0%, CI -0.8% to 3.9%).
- Combining all three asset classes into a single value-weighted private-capital allocation produces an alpha of 2.9% per annum (CI -0.5% to 6.7%), close to but not quite statistically significant at the 5% level.
- Naive alpha estimates — based on standard regressions of private-capital indices on public-market indices, even after using "nowcasted" NAVs to address staleness — remain in the 3.4%-7.2% range, materially overstating the genuine economic alpha. Buyout is the only private asset class that delivers a positive and statistically reliable alpha across all robustness specifications considered.

The full paper is available here.

Summary

Motivation: Institutional investors have substantially increased target allocations to private capital — for example, the aggregate target allocation of U.S. public pension funds has grown from about 8% in 2001 to about 22% in 2021. Estimating the risk-adjusted performance of private capital is therefore central to large-scale capital-allocation decisions. Existing approaches typically estimate the Net Present Value (NPV) of private funds from the perspective of an LP that can freely allocate capital in public markets (Public Market Equivalent, Generalized PME, etc.), but these NPV-style metrics do not translate cleanly into the alpha measures used to compare other asset classes. They also do not reflect important economic realities of private capital investing: NAVs are reported with stale and smoothed valuations, LPs cannot diversify across the full universe of funds because of cost and

monitoring constraints, and LPs cannot control their effective allocation since capital calls and distributions are determined by the GPs. The authors propose a new methodology to estimate alpha at the asset-class level under realistic LP conditions.

Methodology: The data are the MSCI Private Capital (formerly Burgiss) Universe, covering 5,028 U.S. buyout, venture capital, and real estate funds with vintage years 1987-2022, representing $3.64 trillion in fund commitments. The authors compare three alpha estimates: a "naive alpha" using reported NAVs and a regression of value-weighted private-capital indices on public-market equity and bond indices; an alternative naive alpha using the nowcasted NAVs of Brown, Ghysels, and Gredil (2023) to correct for stale valuations; and the authors' "private capital alpha", obtained from realistic LP simulations. In the simulations, an LP starts with a 60/40 public equity/bond allocation in 1987 and 0% in private capital. Each vintage year, the LP commits a fixed fraction of AUM to nine randomly drawn private-capital funds, holding uncalled committed capital in public equities. The annual commitment fraction is calibrated so that the typical LP private-capital allocation matches a target of 20% in the baseline. Capital calls and distributions follow the actual cash flows of the funds drawn, so the realised allocation drifts. After a five-year ramp-up, the LP's portfolio Sharpe ratio is computed and the increase relative to the baseline 60/40 portfolio is converted into an alpha. The simulation is repeated 10,000 times, with median alpha reported and confidence intervals obtained via bootstrap. Robustness specifications vary the NAV-staleness correction, the baseline public portfolio, the target private allocation, the number of funds per year, the ramp-up period, and the liquidity-management asset.

Findings:

- The naive alpha based on reported NAVs is implausibly large (above 6% per annum for all three private asset classes). Using nowcasted NAVs reduces these naive alphas to 3.4%-7.2% — still substantially overstated, because they do not address underdiversification or LP allocation realities.
- The realistic-simulation private-capital alpha is materially lower than naive alphas, highlighting how much of the headline outperformance reflects ignoring the underlying economic frictions.
- Buyout has a private-capital alpha of 2.5% per annum, with a 95% confidence interval from 0.1% to 4.5%. BO is the only private asset class that provides a positive alpha across the entire confidence interval and the only one that reliably remains positive across alternative specifications.
- Venture capital has the largest point estimate (4.5%) but the widest confidence interval (-4.4% to 15.5%). The substantial uncertainty associated with VC investments — driven by extreme dispersion of fund-level outcomes — makes the VC alpha statistically unreliable in this sample.
- Real estate has a smaller positive alpha of 2.0% per annum but a confidence interval from -0.8% to 3.9% that includes zero.
- Combining BO, VC, and RE in a value-weighted private-capital portfolio produces an alpha of 2.9% per annum, with a confidence interval from -0.5% to 6.7%. The combined alpha is close to but not quite statistically distinguishable from zero at conventional levels, but it is clearly an improvement on a public-only baseline portfolio.
- The qualitative messages — buyout reliably positive, VC large but uncertain, RE small and uncertain — survive a wide range of robustness checks varying NAV-staleness

corrections, the public-market baseline, the target private allocation, the number of funds the LP commits to each year, the ramp-up window, and the asset used for liquidity management.

- The simulation framework is flexible and can be adapted to study related questions, including the effect of liquidity shocks on private-capital attractiveness and how changes in commitment pace, fund concentration, or fund-selection skill affect the alpha that LPs can realistically expect to capture.