War-Related News Risk and the Cross Section of Swedish Capital Market Returns

Frans Axelsson Nessim   Ludvig Gustafsson

MSc in Finance — Stockholm School of Economics

Spring 2026

Abstract

We test whether war-related news risk is priced in the cross section of Swedish capital market returns. Using a Boolean keyword search across six major Swedish newspapers, *WarFac* is constructed as the AR(1) innovation in the share of war-related articles over 1993–2024. The factor carries a significant negative risk premium across all factor controls. The premium is concentrated in geopolitical tail events and reversal dynamics are explained through rational and behavioral interpretations. Extending the analysis to fixed income, periods of elevated war news steepen rather than flatten the Swedish yield curve, inconsistent with a flight-to-safety channel. Together, these findings establish that war risk pricing extends beyond U.S. markets, and that a local-language keyword search constitutes a valid and replicable proxy for geopolitical risk.

Keywords:
Geopolitical risk, war risk premium, cross-sectional asset pricing, war hedging, equity markets, fixed income, Swedish capital markets, news-based risk measures

Authors:
Frans Axelsson Nessim (25339)
Ludvig Gustafsson (25335)

Supervisor:
Riccardo Sabbatucci, Associate Professor, Department of Finance, Stockholm School of Economics

Acknowledgments:
We would like to thank our supervisor Riccardo Sabbatucci for his valuable guidance throughout the writing process.

Master Thesis
MSc in Finance
Stockholm School of Economics

I. Introduction

War is the canonical macroeconomic disaster (Barro, 2006), and the risk of it has rarely been more salient. Russia's invasion of Ukraine, ongoing conflict in the Middle East, and rising great-power tensions have returned geopolitical risk to the center of public attention. Yet despite its apparent importance, relatively little is known about how capital markets price war risk in the cross section of stocks and bonds. Theoretical models of rare disaster risk predict that assets co-varying positively with perceived disaster probability provide insurance during tail events and therefore earn lower expected returns in equilibrium (Barro, 2006; Gabaix, 2012; Wachter, 2013). Recent work has also shown that media coverage of war captures this mechanism effectively in U.S. capital markets (Hirshleifer et al., 2025a, 2025b). Whether this finding reflects a universal feature of internationally integrated capital markets or is specific to the U.S. setting remains underexplored, as does whether local-language media can substitute for English-language sources as proxies for global geopolitical risk.

We address both questions, extending the work of Hirshleifer et al. (2025b), using Swedish data. Sweden is a small open economy with a non-reserve currency, exposed to European geopolitical stress, and served by a rich and accessible newspaper archive. Its equity market is well-developed and institutionally transparent, making it a suitable out-of-sample setting for this test: large enough to support cross-sectional asset pricing analysis, yet distinct enough from the U.S. in size, currency, and geopolitical exposure to constitute a meaningful test.

We construct *WarFac* as the AR(1) innovation from a Boolean keyword war share across six major Swedish newspapers and apply the Cochrane (2005, ch. 12) two-pass framework to 1,011 Swedish listed stocks over 1993 to 2024 (innovations in *WarFac* exceeding one standard deviation in absolute value are referred to as *WarFac shocks* throughout). The Swedish war share correlates at $r = 0.684$ with the forward-looking GPR Threats component of Caldara and Iacoviello (2022), validating local-language media as a proxy for global geopolitical risk.

The main finding is a significant negative *WarFac* risk premium of -2.53% per annum ($t_{\text{Shanken}} = -2.00$), robust across all specifications and factor controls. Under a disaster-hedge interpretation, the pricing relationship should manifest during rare geopolitical tail events, where war-hedge stocks outperform war-exposed stocks as their insurance value is realised, rather than as a stable week-by-week return differential. Consistent with this, the unconditional factor-mimicking portfolio (FMP) earns an insignificant 2.40% per annum ($t = 0.40$): a disaster-hedge payoff concentrated in rare tail events cannot survive full-sample time-series averaging. Decomposing returns by shock magnitude reveals the underlying structure directly. The shock-week spread intensifies nonlinearly with shock severity, reaching -181.67% annualized at the 2σ threshold across 3.0% of sample weeks. Post-shock reversals concentrate asymmetrically in the war-hedge portfolio, with a medium-term cumulative abnormal return of -6.61% ($t = -2.09$). Both patterns are consistent with two explanations: rational disaster-hedge pricing and behavioral overpricing of war risk. To test tradability, a war-mimicking portfolio (WMP) is constructed from 92 basis portfolios and earns a significant annualized return of -19.44% ($t = -2.52$), with a weakly significant four-factor alpha of -14.23% ($t = -1.89$). This indicates that the factor's pricing information is embedded in a tradable long-short position, though the result warrants cautious interpretation, partly attributable to the narrow basis set relative to Hirshleifer et al. (2025b). In fixed income, *WarFac* is associated with a decline in bond returns across all specifications, spanning both government and corporate bonds. The negative and significant beta on the government bond spread ($\beta = -0.035$, $t = -2.58$) indicates that long-maturity bonds underperform short-maturity bonds following *WarFac shock* weeks. This pattern is inconsistent with flight-to-safety into Swedish government bonds, suggesting that *WarFac shocks* redirect capital away from non-reserve-currency fixed income markets.

The paper makes three contributions. First, we show that the disaster-hedge pricing mechanism of Hirshleifer et al. (2025b) extends to internationally integrated equity markets outside the U.S., suggesting that war risk pricing is a broad feature of capital markets rather than specific to the U.S. setting. Second, we show that a transparent Boolean keyword search over local-language newspapers constitutes a valid proxy for global geopolitical risk, making cross-country replication feasible without large historical corpora or sophisticated text-processing infrastructure. Third, we introduce a shock-decomposition framework that quantifies the concentration of the war-hedge payoff across tail events of varying severity.

The remainder of the paper is structured as follows. Section II reviews the related literature. Section III describes the data and empirical methodology. Section IV presents and discusses the results. Section V concludes.

II. Related Literature

This paper sits at the intersection of three strands: rare disaster models and their cross-sectional implications, the measurement of geopolitical and war-related risk through text-based proxies, and the international scope of asset pricing anomalies and safe-haven capital flows.

II.A. Rare Disaster Risk, the Equity Premium, and Cross-Sectional Implications

The rare disaster literature originates in the equity premium puzzle of Mehra and Prescott (1985): matching observed risk premia requires implausibly high risk aversion under standard power utility. Rietz (1988) proposes a resolution, where even a small probability of rare consumption disaster substantially increases the required equity premium, and Barro (2006) validates this quantitatively across twentieth-century contractions, showing that a calibrated rare-disaster model can match observed premia with plausible parameters. Subsequent work introduces time variation in disaster risk. Gabaix (2012) develops a variable rare disasters model in which time-varying disaster severity generates cross-sectional variation in expected returns; Wachter (2013) embeds time-varying disaster probability in a representative-agent economy, simultaneously matching the equity premium, excess volatility, and price-dividend ratio predictability. In these models, increases in perceived disaster risk lower asset prices and raise expected future returns.

The cross-sectional prediction is direct: assets co-varying positively with disaster risk earn lower average returns because investors accept lower expected returns in exchange for the hedge. Empirical tests of this prediction have been limited by the difficulty of measuring time-varying disaster risk. Gourio (2008) discusses the cross-sectional implications of rare disaster risk but finds limited empirical support; Hirshleifer et al. (2025b) attribute this null result to the small sample sizes inherent in ex post disaster proxies. Berkman et al. (2011) use 447 international political crises as a proxy and find higher returns for crisis-exposed industries, consistent with the theory, but their approach relies on discrete crisis event counts rather than continuous media discourse.

A closely related cross-sectional prediction arises under a behavioral channel, though through a different mechanism. Under the cumulative prospect theory of Tversky and Kahneman (1992), investors overweight low-probability extreme outcomes, implying that assets providing payoffs in rare disaster states will be overvalued and earn lower expected returns. Daniel et al. (2020) formalize short- and long-horizon behavioral factors arising from investor psychological biases, and provide empirical tools for distinguishing persistent mispricing from rational risk compensation. The two channels generate observationally similar unconditional cross-sectional predictions but differ in their implications for the temporal distribution of the premium: rational disaster-hedge pricing predicts concentration in rare tail events as direct compensation for bearing insurance risk, while the behavioral channel predicts that salience-driven overreaction concentrates around extreme geopolitical episodes and is subsequently reversed as mispricing corrects (Daniel et al., 2020; Tversky & Kahneman, 1992). The relationship between the magnitude of the initial price reaction to war-risk shocks and the extent of any subsequent reversal in the disaster-hedge portfolio therefore provides key statistics for understanding the two channels.

A complementary literature has established that generic downside tail risk is priced in the cross section. Ang et al. (2006) document a premium of approximately 6% per annum for stocks that covary strongly with the market during market declines, a result robust to controls for regular market beta, coskewness, liquidity risk, and size, value, and momentum characteristics. War-related news risk is distinct from this generic downside covariation; Hirshleifer et al. (2025b) demonstrate that war discourse retains significant cross-sectional pricing power after controlling for downside beta and a battery of related tail risk measures, establishing that investors price the particular disaster scenario associated with war over and above their general sensitivity to left-tail market co-movement. However, because this evidence is concentrated in U.S. capital markets, we contribute by providing an out-of-sample test of whether the disaster-hedge pricing mechanism extends beyond the U.S..

II.B. Measuring Geopolitical and War-Related Risk

Measuring time-varying geopolitical risk is empirically challenging, and the literature has moved toward text-based proxies that capture forward-looking risk perception rather than realized events. The template for keyword-based newspaper indices in macroeconomics was established by Baker et al. (2016), whose Economic Policy Uncertainty index measures policy uncertainty from the co-occurrence of economic, uncertainty, and policy terms in newspaper archives, documenting significant effects on investment and employment. Caldara and Iacoviello (2022) develop the GPR index from major international newspapers, decomposed into a threats component (GPR_T, forward-looking) and an acts component (GPR_A, backward-looking), and document significant macroeconomic effects following geopolitical risk spikes. Manela and Moreira (2017) construct NVIX from Wall Street Journal front pages, finding its predictive power for aggregate returns is strongest when driven by war content. Building on these measures, Hirshleifer et al. (2025a) construct a war discourse index from 160 years of New York Times articles using sLDA, showing war discourse predicts aggregate stock and bond returns and outperforms alternative disaster proxies. Hirshleifer et al. (2025b) extend this to the cross section of U.S. equity returns, documenting a significant negative war risk premium and showing that media-discourse-based factors retain incremental pricing power beyond crisis event counts. This resolves the tension left open by Berkman et al. (2011).

A shared feature of these benchmarks is their reliance on English-language sources. Caldara and Iacoviello (2022) acknowledge that their index reflects primarily a North American and British perspective on geopolitical events, noting that it is constructed from newspapers published in the United States, the United Kingdom, and Canada. Whether local-language media in non-Anglophone markets can serve as equally valid proxies for global geopolitical risk perception, and whether the pricing evidence from English-language settings replicates in other information environments, remain open questions in the literature.

A related question is whether the sophisticated text-processing methods employed by these benchmarks are necessary to extract pricing-relevant information from news, or whether simpler keyword-based approaches suffice. Bybee et al. (2023) extract narrative risk factors from Wall Street Journal articles via LDA, IPCA, and group lasso, achieving higher out-of-sample Sharpe ratios than standard factor models; their ICAPM interpretation is that narrative factors proxy for shocks to state variables forecasting future investment opportunities, carrying premia not spanned by traded factors. Bybee et al. (2024) demonstrate that attention to specific news themes forecasts output, employment, and stock returns beyond what standard numerical indicators capture. These papers collectively document that newspaper text contains systematic pricing information not captured by standard factor models, with the pricing channel operating through investors' real-time updating of beliefs about macroeconomic risk. Whether the sophisticated topic models they employ are necessary for this, or whether simpler keyword methods suffice, bears directly on the feasibility of cross-country replication. We address both open questions, testing whether local-language media constitutes a valid proxy for global geopolitical risk and whether a simple keyword approach can substitute for large historical corpora or supervised topic models.

II.C. International Asset Pricing and Safe-Haven Capital Flows

The cross-sectional asset pricing literature has been predominantly focused on U.S. equity markets, with the question of how return predictors translate to international settings remaining an active area of research. Jensen et al. (2023) provide the most comprehensive recent evidence, testing 153 factors across 93 countries and finding a global replication rate broadly comparable to that obtained in the United States alone, a finding that strengthens rather than weakens with the growing number of identified factors. Jacobs and Müller (2020) document, however, that the post-publication return decay which characterizes U.S. anomalies is absent in international markets, a pattern they attribute to structural cross-country barriers to arbitrage capital that effectively segment equity markets. The theoretical foundation for these barriers lies in Shleifer and Vishny (1997), who show that professional arbitrage is conducted by a small number of specialized investors managing outside capital and is subject to performance-based liquidation, rendering it most constrained precisely when prices diverge furthest from fundamental values. Together, these findings indicate that while the same return predictors operate globally, the mechanisms governing their pricing dynamics vary substantially across institutional settings. This institutional heterogeneity is particularly pronounced in the disaster-risk pricing literature: the cross-sectional war risk premium documented by Hirshleifer et al. (2025b) is identified entirely in U.S. data, leaving open whether the mechanism reflects a universal feature of internationally integrated capital markets or is specific to the depth, liquidity, and information environment of the U.S. market.

Whether war risk pricing extends across borders has direct implications for how geopolitical stress propagates through fixed income markets outside the United States, a question that turns critically on the safe-haven status of non-reserve-currency bonds. Baele et al. (2020) provide the most systematic empirical characterization of flight-to-safety episodes, identifying such days across twenty-three countries using only stock and bond returns and a model-averaging approach. Their results document that flight-to-safety episodes are accompanied by outflows from equity funds into government bond and money market funds, and that the effect reflects primarily a flight to quality rather than a flight to liquidity in the corporate bond market. They further find that emerging market equities bear flight-to-safety exposures in excess of their benchmark betas, with the surplus attenuated for more financially integrated markets. The safe-haven status of a bond market depends critically on the reserve currency role of the issuing country: U.S. Treasuries are the primary beneficiary of flight-to-safety flows during global risk episodes, as investors reallocate toward dollar-denominated assets (Hirshleifer et al., 2025a). Krishnamurthy and Vissing-Jorgensen (2012) document the structural basis for this asymmetry, showing that U.S. Treasuries carry a convenience yield averaging 73 basis points per annum reflecting investor demand for their unique safety and liquidity attributes. Given that this demand derives from the dollar's reserve currency status and the institutional role of Treasuries as collateral in financial transactions, it cannot extend to bonds denominated in non-reserve currencies. Whether non-reserve-currency government bonds exhibit safe-haven properties during geopolitical stress remains underexplored, with direct implications for fixed-income hedging strategies outside the U.S. dollar system. We contribute to the literature by investigating whether Swedish government bonds serve as safe havens during geopolitical stress, distinguishing the fixed income response of a peripheral non-reserve-currency market from that of U.S. Treasuries.

III. Data and Methodology

The empirical strategy proceeds in five stages. First, we construct *WarFac* from Swedish newspaper coverage and validate it against established international benchmarks. Second, we test whether *WarFac* is priced in the Swedish equity cross section using the Cochrane (2005, ch. 12) two-pass framework. Third, we sort betas into decile portfolios to investigate the premium's temporal structure. An event study then decomposes the premium into shock and non-shock weeks to identify the pricing mechanism. Fourth, following Hirshleifer et al. (2025b), we construct a WMP and test whether the factor's payoff is accessible through a tradable strategy. Fifth, we investigate whether *WarFac shocks* propagate to Swedish fixed income markets through time-series regressions. Table 7 in Appendix A summarizes all data sources.

III.A. War Share Construction and *WarFac* Estimation

War-related media attention is measured using the Retriever media archive, a comprehensive Swedish media monitoring database covering print and online news. We extract article counts at weekly frequency[1] across six major Swedish outlets: *Dagens Nyheter*, *Svenska Dagbladet*, *Dagens Industri*, *Aftonbladet*, *Expressen*, and *Affärsvärlden*. *Dagens Nyheter* and *Svenska Dagbladet* are Sweden's two largest broadsheet newspapers, *Aftonbladet* and *Expressen* its two highest-circulation evening tabloids, and *Dagens Industri* the leading national business daily (Euromedia Research Group, 2022). *Affärsvärlden* is a weekly magazine read primarily by corporate executives, investors, and financial analysts, and publishes the Affärsvärlden General Index (AFGX), one of Sweden's oldest published equity indices, dating to 1901. Together these six outlets span the broadsheet, tabloid, and financial press, providing broad coverage of the information environment facing Swedish investors throughout the sample period. Three limitations of this construction warrant acknowledgment. First, Retriever is a proprietary archive with an undisclosed indexing methodology, so consistency of coverage across the full sample period cannot be verified. Second, the outlet panel is held fixed throughout, which may underrepresent the information environment facing Swedish investors as the broader media landscape evolved over the sample period. Third, all outlets enter the war share with equal

[1] Hirshleifer et al. (2025a) identify finer-than-monthly estimation as a direction for future research, noting it would increase statistical power. Weekly extraction from the Retriever archive addresses this directly.

weight, abstracting from differences in readership and investor relevance. Nonetheless, the stationarity of the post-1990 war share, confirmed by an Augmented Dickey–Fuller test, provides reassurance that none of these features induces a trend severe enough to invalidate the series. Lastly, since the war share is constructed via a Boolean keyword search, attenuation bias could be present, which a supervised topic modeling approach would avoid. Such noise attenuates *WarFac* betas toward zero in the first pass, biasing the second-pass risk premium toward zero. Any reported premium would therefore likely represent a conservative lower bound on the true magnitude.

For each week, we extract the war-related article count, identified using the Boolean keyword string documented in Appendix B, and the total article count across the same outlets. The Swedish war share is defined as the fraction of total articles classified as war-related:

$$S_t^{\text{war}} = \frac{\text{War articles}_t}{\text{Total articles}_t}. \tag{1}$$

This normalization, which follows Caldara and Iacoviello (2022) and Manela and Moreira (2017), guards against spurious variation driven by changes in total publication volume over time. Coverage prior to 1990 is materially thinner in the archive, and the sample is accordingly restricted to 1990–2024. Figure 1 depicts the war share series and key geopolitical events.

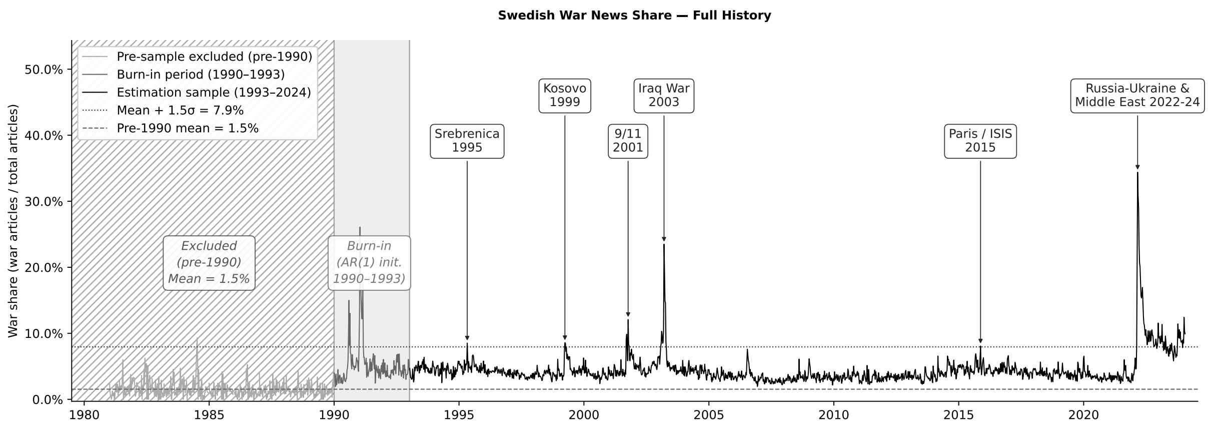

Figure 1. Swedish War News Share. The figure plots the weekly war news share (S_t^{war}) across six major Swedish outlets from the Retriever media archive. Key geopolitical events are annotated. The unshaded region denotes the estimation sample (1993 to 2024); the shaded region provides the burn-in window for the rolling AR(1) used to extract *WarFac* (1990 to 1993).

Because Swedish firms are exposed to global rather than purely domestic geopolitical disruptions, a war share constructed from local-language media is a valid pricing factor only if it captures variation in global conflict risk. We therefore validate the Swedish war share against two established international benchmarks: the GPR index of Caldara and Iacoviello (2022) and the NVIX War index of Manela and Moreira (2017). The GPR index is constructed from automated text searches of ten major international newspapers and decomposes into a forward-looking threats component (GPR_T) and a reactive acts component (GPR_A). The NVIX War index is a war-specific component of the news-implied volatility measure of Manela and Moreira (2017), derived from Wall Street Journal front-page titles and abstracts. Both benchmarks are available only at monthly frequency, so the weekly war share is aggregated to monthly. This is done by computing the ratio of total war article counts to total article counts within each calendar month, preserving the share interpretation of equation (1).

Following Hirshleifer et al. (2025b), we extract *WarFac* as the innovation from an AR(1) fitted to the war share level. The war share exhibits strong serial correlation, and the level series is stationary, implying that AR residuals are well-defined (see Appendix D for full diagnostics).

Since the war share exhibits time-varying dynamics throughout the sample, with multiple episodes visibly shifting the mean and persistence of the series (Figure 1, further reinforced by the rolling stability tests in Appendix D), we fit the AR(1) model over a rolling 156-week window, allowing the parameters to adapt continuously to the prevailing persistence regime.

Specifically, *WarFac* is defined as:

$$S_t^{\text{war}} = \hat{\alpha}_t + \hat{\rho}_t S_{t-1}^{\text{war}} + \hat{\varepsilon}_t, \qquad WarFac_t \equiv \hat{\varepsilon}_t, \tag{2}$$

where $\hat{\alpha}_t$ and $\hat{\rho}_t$ are estimated using only data available through week $t-1$, ensuring *WarFac* is free of look-ahead bias. A persistence bound is applied to the AR(1) slope to ensure meaningful residuals. Details are provided in Appendix C. The residual isolates the unpredictable component of war-related media attention, which Hirshleifer et al. (2025b) argue is the component priced in the cross section of equity returns. The first usable *WarFac* observation falls in January 1993, when 156 weeks of observations first become available.

Sensitivity to the AR specification is assessed via three alternative *WarFac* series for robustness: a rolling two-year window (104 weeks), a rolling five-year window (260 weeks), and a rolling ARMA(1,1). Results are reported in Panel A of Table 12 in Appendix G.

III.B. Swedish Stock Universe

Stock returns and firm characteristics are sourced from the Finbas database, maintained by the Swedish House of Finance Data Center. The dataset covers 2,690 companies and provides weekly data on bid-adjusted prices, book equity, and market capitalization across all Swedish stock market segments.[2] The database retains return histories throughout each stock's trading life, including firms subsequently delisted, acquired, or otherwise removed from the exchange, making it free from survivorship bias. For stocks reported in foreign currencies, prices and market capitalization are converted to SEK using the annual average exchange rate for the relevant calendar year, sourced from Riksbanken.[3] Where a firm has multiple share classes listed simultaneously, returns are aggregated across share classes using a value-weighted approach, to avoid double-counting in portfolio construction and factor estimation. GICS sector classifications,

[2]Four administrative and non-trading categories are excluded throughout, identified by the Finbas codes `SSEUTL`, `SSEA2U`, `EXTERN`, and `INOFF` in line with Aytug et al. (2020).

[3]Using an annual rather than weekly average rate introduces a measurement approximation for firms reporting in foreign currencies. To the extent that exchange rate movements within a year are modest relative to the cross-sectional dispersion in stock returns, this approximation is unlikely to materially affect beta estimates.

used for the sector portfolios in Sections III.C and III.F, are sourced manually for each individual stock from Refinitiv Eikon at the eleven-sector level for the tested stock universe (listed and unlisted stocks) and matched to the Finbas data. The Swedish one-month T-bill rate, sourced from Riksbanken, serves as the risk-free rate.[4] Weekly excess returns are defined as the raw stock return minus this weekly rate.

Stocks for which more than 25% of weekly returns are identically zero are excluded, as sequences of zero returns in thinly traded stocks reflect non-trading and bias beta estimates toward zero. Stocks with fewer than 260 weekly observations are also excluded, ensuring estimates rest on a sufficient time series.[5] After applying these filters, the panel spans 1,011 stocks. Weekly returns are winsorised cross-sectionally at the 1st and 99th percentiles in each week prior to use in any regression.

All pricing tests control for the Carhart (1997) four-factor model, comprising the market excess return (MKT), size (SMB), value (HML), and momentum (UMD). We construct Swedish factors at weekly frequency from Finbas data, similar to the approach of Aytug et al. (2020). The replicated factors are first available in July 1993 due to the applied book equity timing convention, which defines the estimation period used in all equity asset pricing tests as July 1993 to January 2024. The replicated factors are validated against the factors by Aytug et al. (2020), which are published by the Swedish House of Finance and available through December 2019, before being used in any pricing test. Construction details are provided in Appendix E.

III.C. Two-Pass Cross-Sectional Pricing Framework

The primary pricing test follows the two-pass framework of Cochrane (2005, ch. 12). Three nested specifications are estimated: a univariate model with *WarFac* alone, a two-factor model adding the market excess return, and the main specification augmenting the full Carhart (1997) four-factor set with *WarFac*. All tests are conducted at weekly frequency over the full estimation sample.

The 1,011 individual stocks serve as the primary test asset set, maximizing cross-sectional variation in *WarFac* betas.[6] As a robustness check, we re-estimate the two-pass procedure on 92 double-sorted sector portfolios. These are formed by combining 11 equal-weighted GICS sector portfolios with three sets of 27 portfolios double-sorting each sector against size, book-to-market, and momentum terciles respectively, with returns equal-weighted within each cell.[7] Since the Finbas database does not support the full set of accounting characteristics used in Hirshleifer et al. (2025b), the narrower basis set generates less cross-sectional dispersion in *WarFac* betas than in the U.S. setting. Nonetheless, using this complementary test asset set examines whether the pricing relationship survives when the cross-section is structured around economic characteristics

[4] The annualized rate is converted to a weekly decimal as $(r/100)/360 \times 7$ following (Aytug et al., 2020).

[5] Sensitivity to the minimum observation threshold is assessed by reducing it to 156 weeks (three years), which expands the estimation universe to 1,230 stocks.

[6] Cross-sectional regressions are restricted to weeks with at least 30 stocks with valid observations.

[7] Each portfolio subject to a minimum of 5 stocks.

rather than individual firm variation, satisfying the standard of Lewellen et al. (2010) that credible pricing evidence should hold across economically distinct test asset sets.

III.C.1. First Pass: Time-Series Beta Estimation

For each stock i, factor loadings are estimated via full-sample OLS time-series regression:

$$R^e_{i,t} = \alpha_i + \beta_i' \mathbf{F}_t + \eta_{i,t}, \tag{3}$$

where $R^e_{i,t}$ is the weekly excess return of stock i in week t and $\mathbf{F}_t$ is the vector of factors for the relevant specification.

III.C.2. Second Pass: Cross-Sectional Pricing

Time-series mean excess returns are regressed on full-sample betas in a single cross-sectional OLS:

$$\bar{R}^e_i = \lambda_0 + \lambda' \hat{\beta}_i + e_i, \tag{4}$$

where $\bar{R}^e_i$ is the time-series mean excess return of asset i and λ is the vector of risk premia to be estimated. An intercept λ_0 is included to allow for a non-zero zero-beta rate, following Black (1972) and Cochrane (2005, ch. 12). Because first-pass betas are estimated with error, OLS standard errors in the second pass are downward-biased. To handle this, we apply the Shanken (1992) correction, which scales them by the factor:

$$c = 1 + \lambda' \hat{\Sigma}_f^{-1} \lambda, \tag{5}$$

where $\hat{\Sigma}_f$ is the sample covariance matrix of the factors. Reported risk premia $\hat{\lambda}_{WarFac}$ are scaled to a one-standard-deviation *WarFac* innovation throughout.

III.D. The FMP

To complement the cross-sectional pricing evidence with a time-series approach, we construct a FMP on *WarFac* betas. The 1,011 individual stocks are sorted into ten equal-count decile portfolios on their full-sample $\hat{\beta}_{WarFac}$ from the main specification. Decile D1 contains the 102 stocks with the most negative *WarFac* betas, the most war-exposed, while D10 contains those with the most positive betas, the most war-hedged. Returns are equal-weighted within each decile in each week, and the final portfolio is defined as $\text{FMP}_{\text{stock}} = \text{D1} - \text{D10}$.[8]

Sorting stocks on their *WarFac* betas, creating a time-series measure and comparing average returns across deciles allows for further interpretation of the premium distribution across weeks, a distinction the two-pass estimator alone cannot resolve. Further, using the same beta estimates and stock universe as the Cochrane first pass ensures that any return pattern in the decile portfolios

[8] Equal-weighting is consistent with the portfolio construction of Hirshleifer et al. (2025b). In the Swedish market, where capitalization is concentrated among a small number of large-cap names, value-weighting would suppress cross-sectional variation in *WarFac* betas; equal-weighting preserves the full beta dispersion across the size distribution. The value-weighted analogue of the decile spread is also reported as a robustness check.

speaks directly to the cross-sectional pricing relationship identified in the second pass. It also ensures that the FMP and the event study in Section III.E operate on identical portfolios.

FMP alphas are estimated relative to the CAPM, FF3, and FF4 benchmarks using Newey–West standard errors with 26 lags (Newey & West, 1987). Monotonicity of average returns across the sort is assessed via Spearman rank correlation.

III.E. Event Study Design

The full sample FMP captures the average return differential through a time-series approach, but cannot quantify the extent to which the premium is clustered around specific episodes. To address this, we go one step further and conduct an event study that decomposes returns into shock and non-shock weeks to examine whether the *WarFac* premium concentrates in rare geopolitical tail events or is diffusely distributed across all periods.

Abnormal returns (ARs) for each stock-week are computed as the excess return minus the FF4 expected return, evaluated at full-sample betas:

$$\mathrm{AR}_{i,t} = R^e_{i,t} - \left(\hat{\alpha}_i + \hat{\beta}_i^{\mathrm{MKT}} \cdot \mathrm{MKT}_t + \hat{\beta}_i^{\mathrm{SMB}} \cdot \mathrm{SMB}_t + \hat{\beta}_i^{\mathrm{HML}} \cdot \mathrm{HML}_t + \hat{\beta}_i^{\mathrm{UMD}} \cdot \mathrm{UMD}_t\right). \tag{6}$$

WarFac is intentionally excluded from the expected return model so that the spread between portfolio groups captures the full differential response to *WarFac shocks*. Stocks are allocated to the ten decile portfolios constructed in Section III.D, ensuring that the event study and the FMP examine exactly the same portfolios. The equal-weighted D1$-$D10 AR spread serves as the primary test statistic throughout.

All sample weeks are classified into three groups by the absolute magnitude of *WarFac* relative to its standard deviation: shock weeks ($|WarFac_t| >$ threshold), quiet weeks ($0.25\sigma < |WarFac_t| \leq$ threshold), and zero-period weeks ($|WarFac_t| \leq 0.25\sigma$). The D1$-$D10 spread is computed separately within each group and the null of zero mean is tested using bootstrap standard errors with 5,000 draws. Results are reported at three thresholds, 1σ, 1.5σ, and 2σ, to assess sensitivity to the shock definition. The total annual premium decomposes into shock and non-shock contributions as:

$$\lambda_{\text{total}} = \bar{s}_{\text{shock}} \cdot n_{\text{shock}} + \bar{s}_{\text{non-shock}} \cdot n_{\text{non-shock}}, \tag{7}$$

where $\bar{s}$ is the mean weekly D1$-$D10 spread within each group and n is the number of weeks per year in that group.

Further, shock events are examined at the 2σ primary threshold, with 1.5σ used as an alternative for the reversal regression, after applying a 13-week de-clustering window. Where multiple candidate events fall within 13 weeks of each other, only the week with the largest $|WarFac_t|$ is retained, preventing adjacent weeks within the same geopolitical episode from being double-counted. The analysis focuses on positive shock events: war escalations dominate the sample, and the de-clustering procedure leaves too few negative events at either threshold for reliable bootstrap inference. For each identified event, the cumulative abnormal return over a window $[t_1, t_2]$ is defined as:

$$\mathrm{CAR}_i[t_1, t_2] = \sum_{t=t_1}^{t_2} \mathrm{AR}_{i,t}, \tag{8}$$

CARs are computed for D1, D10, and the D1−D10 spread over three windows: contemporaneous $[t-4, t+4]$, short-term $[t+5, t+13]$, and medium-term $[t+5, t+26]$. The latter two windows share the same starting point to facilitate comparison of cumulative drift at different horizons. A reversal regression of the medium-term D1−D10 spread on the contemporaneous spread examines whether larger initial reactions are followed by proportionally larger reversals, providing additional insights into the rational and behavioral channels. Heteroskedasticity-robust (HC3) standard errors are used given the small event samples.

III.F. The WMP

Because *WarFac* is constructed from media coverage, it cannot be traded directly. Following Hirshleifer et al. (2025b), we construct a WMP as a traded proxy for *WarFac shocks*. This tests whether the pricing relationship identified in the two-pass framework corresponds to a payoff accessible to investors. The WMP is constructed by projecting *WarFac* onto the space of traded asset returns using the cross-sectional approach of Lehmann and Modest (1988) as applied by Cooper and Priestley (2011) and implemented in Hirshleifer et al. (2025b). The WMP return in week t is the Fama and MacBeth (1973) cross-sectional slope from regressing portfolio excess returns on their full-sample *WarFac* betas, where these betas are estimated from univariate time-series regressions of portfolio returns on *WarFac* alone:

$$R_t^{\mathrm{WMP}} = \hat{\lambda}_t^{WarFac}, \tag{9}$$

where $\hat{\lambda}_t^{WarFac}$ is estimated using pre-estimated full-sample $\hat{\beta}_{WarFac}$ as the regressor, implicitly going long assets with relatively high *WarFac* betas and short those with relatively low betas, with position sizes scaling with the deviation from the cross-sectional mean.

The WMP requires a reliable cross-sectional slope in every single week, making portfolio aggregation essential: weekly returns at the individual stock level are dominated by idiosyncratic variation, and individual stock betas carry substantially more measurement error than portfolio betas, introducing errors-in-variables attenuation in the weekly slopes (Adrian et al., 2014). Unlike the two-pass estimator, which operates on full-sample means over which idiosyncratic noise averages out, the WMP must identify the factor price from a single week of returns at a time. The basis assets are thus the 92 double-sorted sector portfolios described in Section III.C, which provide cross-sectional variation in *WarFac* betas while preserving the estimation precision that portfolio aggregation affords. Relative to Hirshleifer et al. (2025b), however, this basis set is narrower and spans a smaller portion of the return space, which should generate less dispersion in *WarFac* betas. For that reason, we treat this analysis as a complementary exercise to the main specification.

The WMP is evaluated via its mean return, annualized Sharpe ratio, and spanning tests against the CAPM, FF3, and FF4 benchmarks with Newey–West standard errors using 26 lags (Newey & West, 1987).

III.G. Fixed Income Framework

To complement the equity analysis, we examine whether *WarFac shocks* influence fixed income markets. This is done by estimating full-sample time-series OLS regressions of weekly bond excess returns on *WarFac*. Four fixed income test assets are used: the 2Y and 5Y Swedish government bonds, the 5Y$-$2Y term spread return, and the S&P Sweden Investment Grade Corporate Bond Index.

Swedish government bond yields for the 2Y and 5Y maturities are obtained from Riksbanken at weekly frequency, covering 1990–2024. These are constant-maturity par yields published as part of Riksbanken's official interest rate statistics. Weekly government bond returns are approximated from the par yield series using the standard carry-plus-duration formula:

$$\tilde{R}_t^b = \frac{y_{t-1}}{52} - D_{\text{mod}} \cdot \Delta y_t, \tag{10}$$

where y_{t-1} is the prior-week yield in percent, Δy_t is the change in yield in percentage points, and D_{mod} is the assumed modified duration, set to 2.0 years for the 2Y bond and 5.0 years for the 5Y bond. Equation (10) imposes a fixed modified duration and omits convexity; inferences rest on the sign and significance of the *WarFac* coefficient rather than the precise level of bond returns, and any approximation error is unlikely to be systematically correlated with *WarFac shocks*.

The weekly excess return is defined as:

$$R_t^b = \tilde{R}_t^b - R_{f,t}, \tag{11}$$

where $R_{f,t}$ is the weekly Riksbank one-month T-bill rate. The term spread return is the difference between the 5Y and 2Y excess returns in each week. All government bond excess returns are winsorised at the 1st and 99th percentiles within the sample.

For corporate bonds, we use the S&P Sweden Investment Grade Corporate Bond Index, sourced from S&P Global as a daily total return index and aggregated to weekly frequency by retaining the last available Friday observation. The weekly return is computed as $P_t/P_{t-1} - 1$, with the excess return winsorised at the 1st and 99th percentiles. The index is available from 2016, so results for this asset cover approximately eight years and are reported as subperiod estimates alongside the full-sample government bond findings. Despite its short available time horizon, it helps provide a more complete picture of how war impacts fixed income markets.

For each asset, two specifications are estimated: a univariate regression on *WarFac* alone and an augmented specification controlling for the market excess return, with Newey–West HAC standard errors using 26 lags (Newey & West, 1987) throughout:

$$R_t^b = \alpha^b + \beta_{WarFac}^b \cdot WarFac_t + \beta_{\text{MKT}}^b \cdot \text{MKT}_t + \varepsilon_t^b. \tag{12}$$

Whether the bond market response to *WarFac shocks* is symmetric is assessed by decomposing *WarFac* into its positive and negative components:

$$WarFac_t^+ = \max(WarFac_t, 0), \qquad WarFac_t^- = \min(WarFac_t, 0), \tag{13}$$

and including both components jointly alongside the market control. The null of symmetric response, $H_0: \beta^+ = \beta^-$, is tested using a Wald statistic from the Newey–West covariance matrix. Sensitivity to the lag truncation parameter is assessed at 4 and 13.

IV. Results and Discussion

IV.A. War Share Validation and *WarFac* Properties

IV.A.1. External Validation of the War Share

The Swedish war share co-moves significantly with both the GPR index of Caldara and Iacoviello (2022) and the NVIX War index of Manela and Moreira (2017). Against the GPR index, the Pearson correlation is $r = 0.588$ ($p < 0.001$); against NVIX War, $r = 0.513$ ($p < 0.001$). Lead-lag cross-correlations at lags ± 6 months confirm that the peak correlation with both benchmarks occurs at lag zero. The Swedish series neither systematically leads nor lags its international counterparts, consistent with Swedish media responding to the same global geopolitical events simultaneously rather than with a delay.

In the full sample, the four largest geopolitical episodes are the Gulf War (1990–1991), September 11 and its aftermath (2001), the Iraq War (2003), and the Russian invasion of Ukraine (2022). Each produces a co-elevation above the full-sample mean in both the Swedish war share and the GPR index, highlighting that the Swedish measure responds to the same discrete geopolitical events that drive the established benchmark.

The GPR subindex decomposition proves particularly informative. The Swedish war share correlates more strongly with the forward-looking threats component (GPR_T, $r = 0.684$) than with the reactive acts component (GPR_A, $r = 0.424$). This differential indicates that Swedish media coverage better captures anticipated or escalating war risk rather than merely reporting on ongoing military events. This forward-looking dimension is precisely what Gabaix (2012) and Caldara and Iacoviello (2022) predict should be priced, as investors respond to perceived future disaster probability rather than the realization of conflict itself.

These results collectively establish the Swedish war share as a valid proxy for global geopolitical risk; the subindex decomposition is also informative about the pricing mechanism. The stronger correlation with GPR_T is consistent with rational pricing of anticipated disaster probability (Gabaix, 2012; Wachter, 2013), but does not exclude a behavioral interpretation, since the overweighting of low-probability extreme outcomes under cumulative prospect theory (Tversky & Kahneman, 1992) applies equally to anticipated threats as to realized events. The forward-looking nature of the war share therefore provides partial support for rational pricing of anticipated disaster probability, but is insufficient on its own to distinguish the two channels.

IV.A.2. WarFac Descriptive Statistics and Factor Correlations

The *WarFac* series has an annualized standard deviation of 6.23% and mean of −0.98% ($t = -0.77$), statistically indistinguishable from zero as expected for a residual series. The distribution is strongly right-skewed (skewness = 4.421), a property inherited from the underlying war share since it fluctuates close to zero, leaving substantially more room for large positive innovations than for negative ones. The distribution further exhibits extreme excess kurtosis of 60.468, confirming that *WarFac* is dominated by a small number of discrete geopolitical shock observations, consistent with a factor driven by rare tail events, as examined directly in Section IV.D.

Although near-zero autocorrelation at short lags confirms successful removal of the short-run serial dependence in the war share level series, a Ljung–Box test over ten lags rejects white noise ($p < 0.001$). The rejection is driven by modest positive autocorrelation at lags 3 through 5 that remains small in absolute magnitude. Newey–West standard errors with 26 lags provide valid inference under any remaining serial correlation of this form.

WarFac is essentially orthogonal to the Carhart (1997) four factors, implying that the Shanken correction factor (equation 5) remains close to unity in all specifications, confirming negligible errors-in-variables bias, and that $\hat{\lambda}_{WarFac}$ is insensitive to the inclusion of standard factor controls. Any estimated *WarFac* premium therefore reflects variation in expected returns not spanned by conventional factor exposures. Full descriptive statistics and factor correlations are reported in Appendix D.

IV.B. Two-Pass Cross-Sectional Pricing Tests

IV.B.1. Individual Stocks

Table 1 reports the Cochrane (2005, ch. 12) two-pass cross-sectional regression results for the stocks with valid first-pass betas, with t-statistics applying the Shanken (1992) errors-in-variables correction throughout.

Table 1. Cross-Sectional Pricing of *WarFac*: Individual Stocks

Specification	$\hat{\lambda}_{\text{WarFac}}$ (weekly)	Ann. (%)	t_{Shanken}	c	$Adj.R^2$	N
Univariate	−0.00053**	−2.76	−2.19	1.004	0.38%	1,011
WarFac + MKT	−0.00055**	−2.86	−2.29	1.006	1.18%	1,011
WarFac + FF4 (main)	−0.00049**	−2.53	−2.00	1.020	4.12%	1,011

This table reports Cochrane (2005) full-sample two-pass cross-sectional regression results for the 1,011 Swedish stocks with valid first-pass betas over July 1993 to January 2024. The three specifications are univariate (*WarFac* only), *WarFac*+MKT (adding the market excess return), and the main specification *WarFac*+FF4 (adding the full Carhart (1997) four-factor set). The first pass regresses each stock's weekly excess return on the factor set over the full sample, requiring a minimum of 260 weekly observations. The second pass regresses time-series mean excess returns on full-sample betas in a single cross-sectional OLS, including an intercept following Black (1972). All t-statistics apply the Shanken (1992) errors-in-variables correction and the Shanken correction factor c is reported for each specification. Weekly risk premia are annualized by multiplying by 52. Adjusted R^2 is the cross-sectional adjusted R^2 from the second-pass regression. The full estimated coefficient vector for the main specification is reported in Table 11 in Appendix F. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

In the first pass (equation 3), the full-sample OLS time-series regressions yield a mean R^2 of 0.2% in the univariate specification, rising to 11.5% when the market excess return is added and 14.9% in the main specification augmented by the full Carhart (1997) four-factor set. The cross-section has adequate dispersion in war risk exposure: the standard deviation of *WarFac* betas in the main specification is 0.512, with a range of $[-3.059, +2.855]$, providing sufficient variation to identify a risk price in the second pass.

The second-pass results (equation 4) indicate that *WarFac* is priced. The estimated risk premium is negative and statistically significant at the 5% level across all three specifications. In the univariate specification the annualized premium is −2.76% ($t_{\text{Shanken}} = -2.19$); adding the market excess return yields −2.86% ($t_{\text{Shanken}} = -2.29$); and in the main specification, where *WarFac* augments the full four-factor set, the premium is −2.53% per annum ($t_{\text{Shanken}} = -2.00$). The variation across specifications is less than 0.35 percentage points, reflecting the near-zero factor correlations documented in Section IV.A.2. The Shanken correction factor of $c = 1.020$ confirms that errors-in-variables bias is negligible. The cross-sectional R^2 of 4.1% is low but standard for individual-stock cross-sections, where idiosyncratic noise limits explanatory power.

The baseline premium of −2.53% per annum is negative and statistically significant across all alternative *WarFac* constructions. A two-year rolling window yields −2.89% ($t_{\text{Shanken}} = -2.28$), a five-year rolling window yields −2.14% ($t_{\text{Shanken}} = -1.70$), and the two alternatives correlate with the baseline at $r = 0.990$ and $r = 0.989$ respectively, confirming that window length does not materially affect the informational content of *WarFac*. A rolling ARMA(1,1) specification yields a stronger estimate of −7.21% ($t_{\text{Shanken}} = -3.24$, significant at the 1% level), though its low correlation with the baseline ($r = 0.583$) indicates a materially different filter than the AR(1) derived from Hirshleifer et al. (2025b). The result nonetheless confirms that alternative filtering approaches to the same underlying news series produce negative and significant war risk premia.

Further, reducing the minimum observation threshold to 156 weeks, which increases the estimation universe to 1,230 stocks, yields -4.45% ($t_{\text{Shanken}} = -3.99$), consistent with the baseline and significant at the 1% level. As a cross-sectional check, the Cochrane two-pass is applied to 92 double-sorted sector portfolios which yields -14.42% per annum ($t_{\text{Shanken}} = -1.91$, significant at the 10% level). Full robustness results are reported in Table 12 in Appendix G.

The results establish that stocks with higher positive *WarFac* betas earn lower average returns, directionally consistent with the U.S. evidence of Hirshleifer et al. (2025b) and providing out-of-sample support for the disaster-hedge mechanism in a distinct capital market setting. The cross-sectional risk-return tradeoff slopes downward in the *WarFac* beta dimension, with war-hedge stocks earning lower expected returns and war-exposed stocks earning higher expected returns. Two interpretations of this downward slope are consistent with the evidence. Under the rational disaster-hedge interpretation (Barro, 2006; Gourio, 2008; Wachter, 2013), an increase in war discourse raises investor perceptions of rare macroeconomic disaster probability, and assets that perform well during these episodes provide valuable insurance, earning lower expected returns as compensation. Under the behavioral interpretation (Daniel et al., 2020; Tversky & Kahneman, 1992), investors overweight the probability of rare war-related disasters, bidding up prices of disaster-hedge stocks beyond what fundamentals justify. While the two-pass framework provides limited evidence for separating the two channels, the FMP and event study in Sections IV.C and IV.D examine the temporal structure of the premium and bear more directly on the two candidate interpretations.

IV.C. The FMP

The FMP provides a time-series complement to the cross-sectional evidence: unlike the Cochrane second pass, which operates on full-sample mean returns, the FMP is sensitive to how the premium is distributed over time. Table 2 reports descriptive statistics for the extreme equal-count decile portfolios sorted on full-sample $\hat{\beta}_{WarFac}$ and the corresponding FMP.

Table 2. *WarFac*: The FMP

Portfolio	Mean $\hat{\beta}_{WarFac}$	Ann. (%)	t
D1 (lowest β)	-0.901	11.05	1.96**
D10 (highest β)	$+0.985$	8.65	1.33
$\text{FMP}_{\text{stock}}$ (D1$-$D10)	-1.843	2.40	0.40

This table reports descriptive statistics for the extreme decile portfolios of stocks sorted on full-sample $\hat{\beta}_{WarFac}$ from the main specification (*WarFac*+FF4), estimated over July 1993 to January 2024. D1 contains the 102 stocks with the most negative *WarFac* betas (most war-exposed); D10 contains the stocks with the most positive *WarFac* betas (war-hedges). Returns are equal-weighted within each decile in each week. Ann. (%) is the annualized mean excess return. t is the t-statistic for the null that mean excess return equals zero. $\text{FMP}_{\text{stock}}(\text{D1} - \text{D10})$ is the factor-mimicking portfolio. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

The individual stock decile portfolios reveal a directionally consistent pattern. D1, containing the 102 stocks with the most negative *WarFac* betas (mean $\hat{\beta} = -0.901$), earns 11.05% per annum, while D10, containing the stocks with the most positive betas (mean $\hat{\beta} = +0.985$), earns 8.65% per annum. The direction of this spread is consistent with a negative *WarFac* risk premium: stocks bearing greater war risk exposure earn higher average returns.

The stock-level FMP ($\text{FMP}_{\text{stock}} = \text{D1} - \text{D10}$) earns a raw return of 2.40% per annum ($t = 0.40$) and a FF4 alpha of 2.09% per annum ($t = 0.42$), neither statistically significant. The portfolio loads strongly negatively on *WarFac* ($\hat{\beta}_{WarFac} = -1.843$, as expected by construction) and is approximately market-neutral ($\hat{\beta}_{\text{MKT}} = 0.024$). The Spearman rank correlation between decile number and average return is -0.297 ($p = 0.405$), negative and directionally consistent with the two-pass result but not statistically significant.

The unconditional insignificance of $\text{FMP}_{\text{stock}}$ is the natural result under a tail-concentrated premium: rare shock-week signals are overwhelmed in full-sample averages by the positive spread that characterizes quiet periods. The value-weighted decile specification yields a spread of 6.75% ($t = 0.80$), confirming that the attenuation is not mechanically driven by equal-weighting of small-cap names. The following event study in Section IV.D quantifies this structure directly.

IV.D. Event Study: Mechanism and Interpretation

The two-pass estimator finds a significant negative *WarFac* premium and the FMP shows it does not manifest as a stable week-by-week spread. These results are jointly consistent with the disaster-hedge payoff concentrating in rare geopolitical tail events rather than diffusing across the full sample. Concentration in tail events is a necessary feature of rational disaster-hedge pricing but does not by itself exclude the behavioral channel of concentrated overreaction to salient extreme episodes (Daniel et al., 2020; Tversky & Kahneman, 1992). We further investigate the shock concentration of the premium in Section IV.D.1 and cumulative abnormal returns around identified shock events in Section IV.D.2.

IV.D.1. Shock Concentration of the Premium

Table 3 reports the mean D1 minus D10 abnormal return spread within each week classification bucket at three shock thresholds.

Table 3. Conditional D1−D10 Spread by *WarFac* Magnitude

Threshold	Period type	N weeks	Mean (%/wk)	Ann. (%)	t-stat (bootstrap)
Unconditional					
	All weeks	1,242	+0.120	+6.25	+1.06
Panel A: Shock threshold $\lvert WarFac \rvert > 1.0\sigma$					
1.0σ	Shock ($\lvert WF \rvert > 1.0\sigma$)	216	−0.525	−27.30	−1.27
1.0σ	Quiet (0.25σ–1.0σ)	690	+0.287	+14.93	+2.11**
1.0σ	Zero period ($< 0.25\sigma$)	336	+0.192	+10.01	+1.07
Panel B: Shock threshold $\lvert WarFac \rvert > 1.5\sigma$					
1.5σ	Shock ($\lvert WF \rvert > 1.5\sigma$)	88	−1.527	−79.40	−2.07**
1.5σ	Quiet (0.25σ–1.5σ)	818	+0.268	+13.92	+2.00**
1.5σ	Zero period ($< 0.25\sigma$)	336	+0.192	+10.01	+1.06
Panel C: Shock threshold $\lvert WarFac \rvert > 2.0\sigma$					
2.0σ	Shock ($\lvert WF \rvert > 2.0\sigma$)	37	−3.494	−181.67	−2.54**
2.0σ	Quiet (0.25σ–2.0σ)	869	+0.246	+12.80	+1.86*
2.0σ	Zero period ($< 0.25\sigma$)	336	+0.192	+10.01	+1.07

D1−D10 spread in equal-weighted FF4-adjusted abnormal returns. D1 contains the 102 stocks with the most negative *WarFac* betas (most war-exposed); D10 contains the stocks with the most positive *WarFac* betas (war-hedges). N weeks denotes the number of weeks with valid spread observations within each bucket. Bootstrap standard errors (5,000 draws). std($WarFac$) $= 0.00864$. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

The results display a pronounced threshold effect. At the 1σ threshold (216 shock weeks with valid spread observations, 17.4% of the estimation sample), the mean weekly D1 minus D10 spread during shock weeks is −27.30% annualized ($t = -1.27$), statistically indistinguishable from zero, while quiet weeks show a positive and significant spread and zero-period weeks show a positive but insignificant spread. The picture changes markedly at the 1.5σ threshold: the shock-week spread reaches −79.40% annualized ($t = -2.07$, significant at the 5% level). At the 2σ threshold, with just 3.0% of weeks classified as shocks, the shock-week spread intensifies to −181.67% annualized ($t = -2.54$, significant at the 5% level). This nonlinear intensification as the threshold rises distinguishes a genuine pricing relationship from a statistical artefact: a spurious result driven by noise would attenuate as the shock sample shrinks, not intensify.

The shock decomposition of equation (7) provides the central quantitative finding. At the 2σ threshold, shock weeks contribute −5.41% per year while non-shock weeks contribute +11.67% per year, netting to the unconditional spread of +6.26% per annum. The opposing signs reflect two distinct forces: during rare geopolitical tail events, war-hedge stocks outperform war-exposed stocks as their insurance value is realised; during the remaining 97% of weeks, war-exposed stocks earn a positive spread as compensation for bearing disaster risk in normal times. The total unconditional premium is therefore the net of these two components. Figure 2 illustrates this concentration visually across all three shock thresholds.

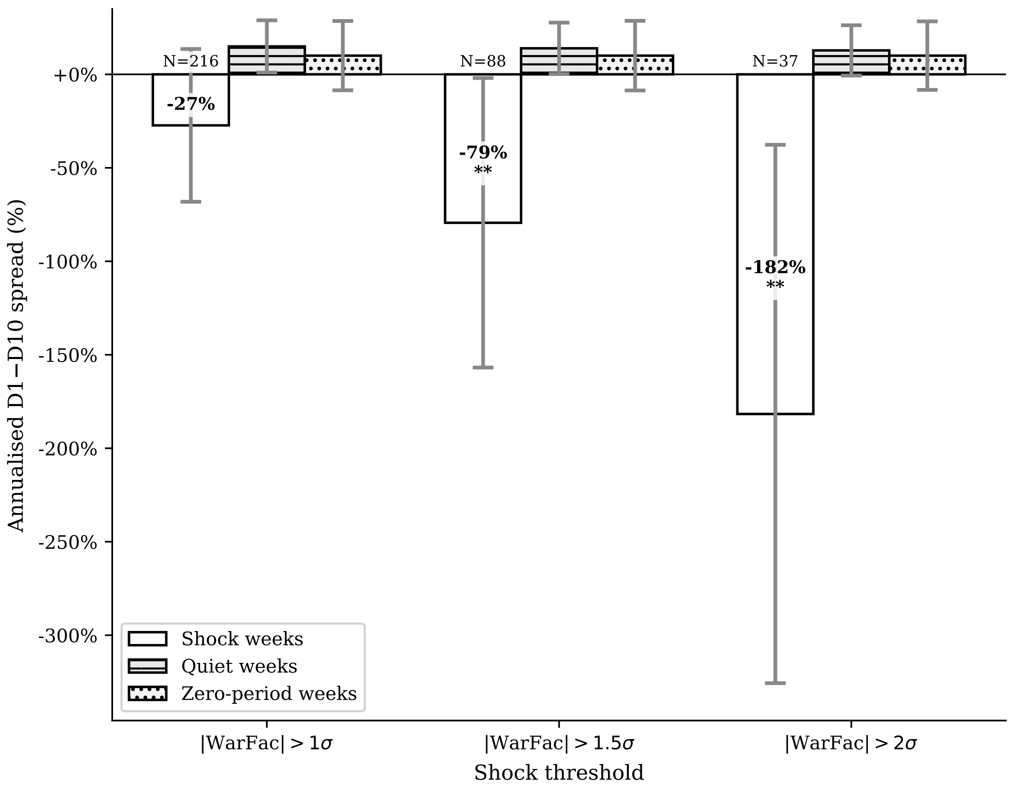

Figure 2. Conditional D1 minus D10 Spread at Three Shock Thresholds. This figure plots the annualized mean D1 minus D10 abnormal return spread separately for shock, quiet, and zero-period weeks at the 1σ, 1.5σ, and 2σ thresholds. Shock weeks are defined as $|WarFac| >$ threshold. Error bars represent 95% confidence intervals ($\pm 1.96\,\mathrm{SE}$) from bootstrap standard errors (5,000 draws). N denotes the number of shock weeks with valid spread observations at each threshold.

This decomposition connects the two estimators. The Cochrane second pass operates on full-sample mean returns, which already incorporate the shock-week contribution; the rare large negative signal survives in the cross-sectional mean even when it cannot survive sequential time-series averaging. The two-pass premium of -2.53% per annum and the shock-week spread of -181.67% annualized are two windows onto the same underlying pricing relationship.

The shock-concentration structure is consistent with rational disaster-hedge pricing (Barro, 2006; Gabaix, 2012; Gourio, 2008; Wachter, 2013), under which disaster-hedge stocks earn lower unconditional returns as equilibrium compensation for the insurance they provide against geopolitical tail events. A simple behavioral channel, under which mispricing unwinds gradually through arbitrage to generate a persistent and broadly distributed return differential, is inconsistent with the positive quiet-week spread and the nonlinear intensification of the shock-week signal across thresholds. The concentrated overreaction framework (Daniel et al., 2020; Tversky & Kahneman, 1992), which represents a more sophisticated behavioral interpretation, is nonetheless not excluded by the shock-concentration evidence. The post-shock dynamics examined in the following section IV.D.2 provide additional insights into both channels.

IV.D.2. Cumulative Abnormal Returns Around Shock Events

At the 2σ threshold, after applying the 13-week de-clustering window, we identify 18 positive shock events. The identified events correspond to major recognized geopolitical episodes, confirming that *WarFac* captures genuine real-time geopolitical tail risk; the full event list is reported in Appendix H. Table 4 reports cumulative abnormal returns around these events (one excluded for insufficient post-event data, $N = 17$).

Table 4. Cumulative Abnormal Returns Around Positive Shock Events (2σ, $N = 17$)

Window	Portfolio	Mean CAR	Bootstrap SE	t-stat
Panel A: Contemporaneous $[t-4, t+4]$				
$[t-4, t+4]$	D1 (most war-exposed)	−0.77%	2.04%	−0.38
$[t-4, t+4]$	D10 (war-hedge)	+8.57%	4.65%	+1.84*
$[t-4, t+4]$	D1−D10	−9.34%	3.72%	−2.51**
Panel B: Short-term $[t+5, t+13]$				
$[t+5, t+13]$	D1	+1.09%	2.36%	+0.46
$[t+5, t+13]$	D10	−3.62%	2.09%	−1.74*
$[t+5, t+13]$	D1−D10	+4.71%	2.19%	+2.15**
Panel C: Medium-term $[t+5, t+26]$				
$[t+5, t+26]$	D1	+0.63%	3.43%	+0.18
$[t+5, t+26]$	D10	−6.61%	3.16%	−2.09**
$[t+5, t+26]$	D1−D10	+7.24%	3.40%	+2.13**

CARs are sums of weekly FF4-adjusted abnormal returns over the indicated window, averaged across 17 positive shock events (18 positive events identified at the 2σ threshold after applying a 13-week de-clustering window; one event excluded for insufficient post-event data). D1 contains the 102 stocks with the most negative *WarFac* betas (most war-exposed); D10 contains the stocks with the most positive *WarFac* betas (war-hedges). Bootstrap standard errors (5,000 draws). ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

The contemporaneous D1 minus D10 spread of −9.34% ($t = -2.51$, significant at the 5% level) confirms that during geopolitical shock events, war-hedge stocks significantly outperform war-exposed stocks, consistent with the disaster-hedge payoff identified in Table 3. During positive *WarFac shock* episodes, war-hedge stocks (D10) earn a contemporaneous cumulative abnormal return of +8.57% ($t = +1.84$, weakly significant at the 10% level), indicating their positive betas drive co-movement with war news. War-exposed stocks (D1) are essentially unchanged (CAR = −0.77%, insignificant), consistent with the interpretation that the war-risk premium on D1 stocks represents compensation earned over the full cycle rather than through contemporaneous underperformance during escalation episodes.

In the weeks following shock events, the D1 minus D10 spread reverses to +4.71% in the short-term window ($t = +2.15$, significant at the 5% level) and remains significantly positive at +7.24% in the medium-term window ($t = +2.13$, significant at the 5% level). The reversal is driven entirely by D10, with cumulative abnormal returns of −3.62% ($t = -1.74$) and −6.61%

($t = -2.09$) in the short and medium-term windows respectively, while D1 remains flat throughout. This asymmetric pattern is consistent with both rational mean-reversion of perceived disaster probability under Wachter (2013) and Gabaix (2012) and transient behavioral overpricing of disaster insurance during extreme geopolitical episodes (Daniel et al., 2020; Tversky & Kahneman, 1992). Figure 3 plots the CAR paths for D1, D10, and the D1 minus D10 spread across the three event windows.

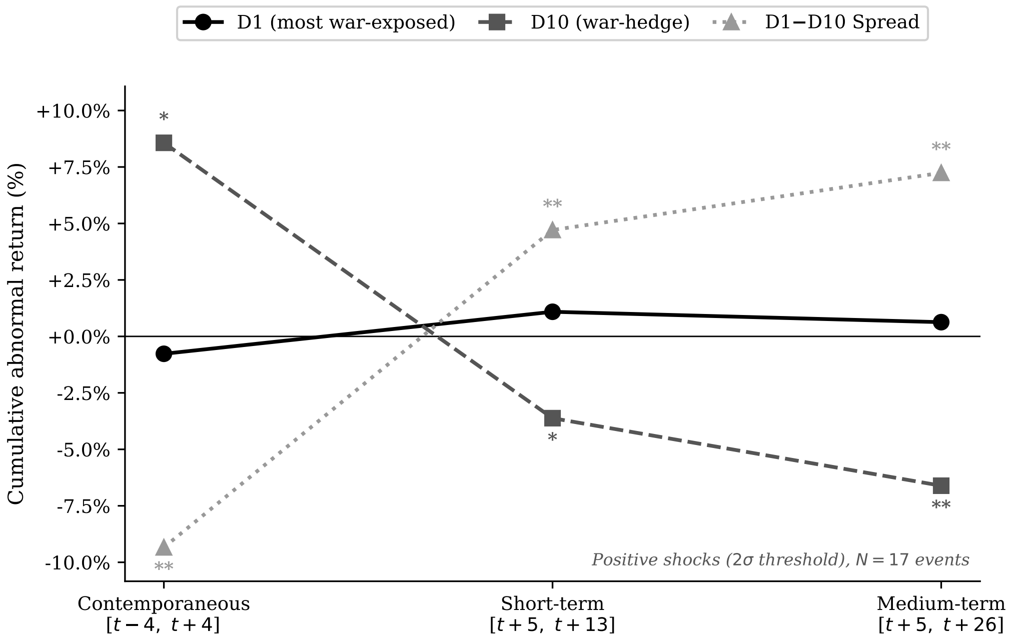

Figure 3. Cumulative Abnormal Returns Around Positive Shock Events (2σ, $N = 17$). This figure plots mean cumulative abnormal returns for D1 (war-exposed stocks, most negative *WarFac* beta), D10 (war-hedge stocks, most positive *WarFac* beta), and the D1 minus D10 spread across three event windows: contemporaneous $[t-4, t+4]$, short-term $[t+5, t+13]$, and medium-term $[t+5, t+26]$, where the latter two windows share the same starting point to facilitate comparison of cumulative drift at different horizons.

To probe the mechanism further, we estimate a reversal regression in which the medium-term post-shock D1 minus D10 spread is regressed on the contemporaneous spread across the $N = 17$ identified events, providing a direct test of whether larger initial reactions are systematically followed by proportionally larger subsequent reversals. The estimated slope is $b = -0.363$ ($t = -1.67$) at the 2σ primary threshold, marginally significant at the 10% level and consistent with transient overpricing. At the 1.5σ alternative threshold, where the event sample expands to $N = 27$, the slope reaches $b = -0.443$ ($t = -2.23$), significant at the 5% level. At the 1σ threshold ($N = 35$), the slope is $b = -0.276$ ($t = -1.24$), negative but insignificant.[9] The slope, which achieves statistical significance at the primary shock thresholds, is consistent with

[9]The non-monotonic pattern of significance across thresholds, stronger at 1.5σ than at either 2σ or 1σ, reflects two opposing forces: lowering the threshold increases the event sample and statistical power, but also introduces weaker shocks with smaller and noisier contemporaneous reactions that attenuate the slope estimate. The 1.5σ threshold represents the point at which the power gain dominates; at 1σ the attenuation from marginal events prevails.

both channels. Under rational mean-reversion of disaster probability, larger shocks generate proportionally larger initial and subsequent price movements; under behavioral overreaction, larger initial overreactions correct proportionally as salience fades. The reversal regression therefore does not discriminate between the two interpretations.

The evidence collectively establishes that *WarFac* is a priced source of systematic variation in expected returns, with the disaster-hedge payoff concentrated in rare geopolitical tail events. The shock-concentration structure rules out a simple diffuse behavioral channel but is consistent with both rational disaster-hedge pricing and the concentrated overreaction framework (Daniel et al., 2020; Tversky & Kahneman, 1992). As Hirshleifer et al. (2025b) note, both rational and behavioral channels predict that expected returns decrease with *WarFac* loadings; this result extends the disaster-hedge framework of Gourio (2008) as documented for U.S. equity markets to the Swedish setting. The forward-looking nature of the war share, as documented in Section IV.A.1, provides partial support for rational pricing of anticipated disaster probability, though neither channel can be cleanly excluded by the available evidence.

IV.E. The WMP

The evidence thus far establishes that the *WarFac* risk premium is priced in the cross section and that the disaster-hedge payoff concentrates in geopolitical tail events. We now test whether it is accessible to investors in the Swedish market. Table 5 reports summary statistics and spanning tests.

Table 5. WMP: Summary Statistics and Spanning Tests

	Mean return		Spanning test alpha		
	Ann. (%)	t	CAPM	FF3	FF4
WMP	−19.44	−2.52**	−20.60***	−20.18***	−14.23*

This table reports summary statistics and spanning tests for the WMP constructed by projecting *WarFac* onto the space of traded asset excess returns using the 92 double-sorted sector portfolios described in Section III.F. Mean return is the annualized mean weekly WMP return. Spanning test alphas are annualized intercepts from time-series regressions of the WMP return on the CAPM, Fama and French (1993) three-factor (FF3), and Carhart (1997) four-factor (FF4) models, with Newey–West standard errors using 26 lags (Newey & West, 1987). The sample covers July 1993 to January 2024. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

The WMP results are directional but constrained by the narrow basis set. The mean return of −19.44% per annum is significant at the 5% level, with an annualized absolute Sharpe ratio of 0.52, comparable to the equity market premium over the same period but substantially below the 1.73 reported by Hirshleifer et al. (2025b) for the U.S. sample. The FF4 alpha of −14.23% ($t = -1.89$) is marginally significant at the 10% level. The CAPM and FF3 alphas are significant at the 1% level, suggesting the return is not fully explained by the market, size, or value factors, but the FF4 result indicates momentum absorbs much of the residual variation. This is economically plausible: during periods of elevated war risk, investors accumulate disaster-hedge stocks, which consequently appear as recent winners in the momentum ranking. Controlling for momentum therefore partially subsumes the war risk signal, as the two factors capture overlapping variation during geopolitical escalation episodes. The basis set of only 92 portfolios is far narrower than the U.S. setting of Hirshleifer et al. (2025b) and spans a smaller portion of the return space, constraining cross-sectional dispersion in *WarFac* betas and making precise mimicking portfolio construction harder than in the U.S. setting. The WMP is therefore a directional complement to the cross-sectional evidence rather than an independent confirmation of a tradable premium. That significant CAPM and FF3 alphas, as well as a significant mean return, nonetheless emerge despite this constraint is noteworthy.

IV.F. Evidence From the Swedish Fixed Income Market

IV.F.1. Main Results

The preceding results indicate that *WarFac shocks* are priced in the Swedish equity cross-section. We next examine whether this pricing extends to fixed income and whether Swedish government bonds serve as safe havens during geopolitical stress, as Hirshleifer et al. (2025a) document for U.S. Treasuries.

Table 6 reports full-sample time-series OLS regression results for the four fixed income test assets (equation 12). Newey–West standard errors with 26 lags are used throughout; sensitivity to the lag choice is assessed in Appendix I.

Table 6. Bond Market Response to *WarFac* Shocks

Asset	$\hat{\beta}_{WarFac}$	Ann. (%)	t_{NW}
Panel A: Government bonds			
2Y Government Bond	−0.00842	−43.79	−1.08
5Y Government Bond	−0.04305**	−223.84**	−2.09**
Panel B: Term structure			
5Y−2Y Spread	−0.03463***	−180.05***	−2.58***
Panel C: Corporate bonds			
IG Corporate Bond	−0.01317	−68.49	−0.60

This table reports full-sample time-series OLS regressions of weekly fixed income excess returns on *WarFac* and the market excess return, with an intercept included. Ann. (%) annualizes $\hat{\beta}_{WarFac}$ by multiplying by 52. t-statistics use Newey–West HAC standard errors with 26 lags (Newey & West, 1987). Government bond excess returns are constructed from Riksbanken constant-maturity par yields using the duration approximation in Section III.G. The term spread return is the 5Y minus 2Y excess return each week. The government bond sample covers January 1993 to January 2024. The IG Corporate Bond return is the weekly excess return on the S&P Sweden Investment Grade Corporate Bond Index, available from April 2016 and reported as a subperiod estimate. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

The five-year government bond displays a negative *WarFac* coefficient statistically significant at the 5% level ($t = -2.09$). A one-standard-deviation *WarFac shock* is associated with a decline of approximately 1.93% per annum in the 5Y bond excess return, providing an economically interpretable magnitude alongside the large raw coefficient. The 5Y minus 2Y term spread return shows a strongly significant negative coefficient at the 1% level ($t = -2.58$), while the two-year bond coefficient is negative but not statistically significant ($t = -1.08$). The concentration of significance at the five-year maturity is consistent with duration amplifying sensitivity to long-run capital outflow and risk repricing at the long end of the yield curve, and the term spread result captures this yield curve steepening directly.

The corporate bond coefficient is negative across all specifications ($t = -0.60$ in the augmented specification) but not statistically significant, consistent with the limited power of an approximately eight-year sample. Baele et al. (2020) document that flight-to-safety episodes reflect primarily a flight to quality, rotating capital from corporate into government bonds. The directional consistency across both asset classes, though inconclusive given the short corporate bond sample, is more consistent with a broad capital outflow from Swedish fixed income than with reallocation within it.

The Wald test for asymmetric response (equation 13) does not reject symmetry for any of the four test assets, and we report the symmetric specification as the main result throughout. The five-year bond result is statistically significant at the 5% level across all three lag specifications: $t = -1.99$ at NW = 4, $t = -2.06$ at NW = 13, and $t = -2.09$ at NW = 26. The term spread coefficient is statistically significant at the 5% level for NW = 4 and NW = 13 ($t = -2.42$ and -2.54 respectively) and at the 1% level for NW = 26 ($t = -2.58$). The two-year bond and corporate bond coefficients are negative but not statistically significant across all lag choices. The point estimate $\hat{\beta}_{WarFac}$ is identical across lag choices for each asset by construction; only the standard error varies. Table 15 in Appendix I reports robustness tests for the main bond results.

These results indicate that *WarFac shocks* are associated with rising long-term Swedish government yields, inconsistent with a flight-to-safety channel. Unlike U.S. Treasuries, for which Hirshleifer et al. (2025a) document a flight-to-quality channel, Swedish government bonds at five-year maturity move adversely with war risk, consistent with the absence of reserve-currency safe-haven status for SEK-denominated bonds. The broad negative sign across all four fixed income test assets indicates a generalized capital outflow from Swedish markets during geopolitical stress.

This outflow channel is consistent with the equity pricing result. If global investors treat Sweden as a peripheral market to exit during disaster states, Swedish stocks that bear geopolitical risk are particularly valuable as return-generating assets, consistent with the war risk premium documented for the U.S. equity markets also persisting in a smaller and less liquid setting. For investors seeking geopolitical portfolio protection in Sweden, these results suggest that equity-based strategies tilting toward war-hedged stocks may prove more reliable than long-duration government bond allocations.

V. CONCLUSION

This paper asks whether war-related news risk is priced in the cross section of Swedish capital market returns and whether local-language newspaper coverage constitutes a valid proxy for global geopolitical risk. Applying the Cochrane (2005, ch. 12) two-pass framework to 1,011 Swedish listed stocks over 1993 to 2024, we find a significant negative *WarFac* risk premium of -2.53% per annum, robust across all specifications and factor controls. The disaster-hedge payoff concentrates sharply in geopolitical tail events rather than diffusing across the full sample, ruling out a simple behavioral channel while remaining consistent with both rational disaster-hedge pricing (Barro, 2006; Gabaix, 2012; Gourio, 2008; Wachter, 2013) and the concentrated overreaction framework (Daniel et al., 2020; Tversky & Kahneman, 1992). A WMP on 92 double-sorted sector portfolios earns a significant negative return, best interpreted as directional support for the cross-sectional evidence given the constraints of the narrow basis set. In fixed income, *WarFac* shocks are associated with rising rather than falling long-term Swedish government yields, pointing to capital flows toward reserve-currency safe havens rather than domestic fixed income during geopolitical stress.

These findings establish that the disaster-hedge pricing mechanism of Hirshleifer et al. (2025b) extends beyond U.S. capital markets, and that a simple Boolean keyword search over local-language newspapers constitutes a valid proxy for global geopolitical risk, making cross-country replication feasible without large historical corpora or more sophisticated models. We also propose a shock-week decomposition of the premium that reveals the underlying pricing mechanism of the factor.

The most important question left to explore is a more granular separation of rational and behavioral channels. One suggestion for future research would be to examine firm-level cash flow realizations around *WarFac shocks*. If the premium reflects genuine disaster compensation, war-exposed assets should exhibit systematically different cash flow dynamics during high-*WarFac* periods, whereas a behavioral explanation would predict price movements decoupled from fundamentals. Complementary evidence could come from ownership data, where the response of retail versus institutional investors around shock events may speak to the role of sentiment versus risk-based demand. Further, expanding the analysis to other Nordic markets using the same Boolean keyword approach would broaden the evidence base for text-based war risk measurement without recourse to supervised topic modeling. It would also enable construction of a richer set of characteristic-sorted portfolios following Hirshleifer et al. (2025b), generating greater cross-sectional dispersion in *WarFac* betas and providing a more powerful test of the pricing relationship. Finally, expanding the sample beyond January 2024 would provide out-of-sample evidence during a period of continued geopolitical disruption, where persistently elevated war coverage may shift the underlying dynamics of the pricing relationship identified.

The evidence presented here supports the view that rare disasters, and wars in particular, are important for understanding the cross section of expected returns on capital. In a world of heightened geopolitical uncertainty, understanding how capital markets price that risk is of direct relevance both to investors constructing robust portfolios and to researchers seeking to explain the persistent heterogeneity in expected returns across firms.

REFERENCES

Adrian, T., Etula, E., & Muir, T. (2014). Financial intermediaries and the cross-section of asset returns. *Journal of Finance*, *69*(6), 2557–2596. https://doi.org/10.1111/jofi.12189

Ang, A., Chen, J., & Xing, Y. (2006). Downside risk. *Review of Financial Studies*, *19*(4), 1191–1239. https://doi.org/10.1093/rfs/hhj035

Aytug, H., Fu, Y., & Sodini, P. (2020). Construction of the Fama–French–Carhart four factors model for the Swedish stock market using the Finbas data [Technical report]. Swedish House of Finance Data Center.

Baele, L., Bekaert, G., Inghelbrecht, K., & Wei, M. (2020). Flights to safety. *Review of Financial Studies*, *33*(2), 689–746. https://doi.org/10.1093/rfs/hhz055

Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, *131*(4), 1593–1636. https://doi.org/10.1093/qje/qjw024

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, *121*(3), 823–866. https://doi.org/10.1162/qjec.121.3.823

Berkman, H., Jacobsen, B., & Lee, J. B. (2011). Time-varying rare disaster risk and stock returns. *Journal of Financial Economics*, *101*(2), 313–332. https://doi.org/10.1016/j.jfineco.2011.02.019

Black, F. (1972). Capital market equilibrium with restricted borrowing. *Journal of Business*, *45*(3), 444–455. https://doi.org/10.1086/295472

Bybee, L., Kelly, B., Manela, A., & Xiu, D. (2024). Business news and business cycles. *Journal of Finance*, *79*(5), 3105–3147. https://doi.org/10.1111/jofi.13377

Bybee, L., Kelly, B., & Su, Y. (2023). Narrative asset pricing: Interpretable systematic risk factors from news text. *Review of Financial Studies*, *36*(12), 4759–4787. https://doi.org/10.1093/rfs/hhad042

Caldara, D., & Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, *112*(4), 1194–1225. https://doi.org/10.1257/aer.20191823

Carhart, M. M. (1997). On persistence in mutual fund performance. *Journal of Finance*, *52*(1), 57–82. https://doi.org/10.1111/j.1540-6261.1997.tb03808.x

Cochrane, J. H. (2005). *Asset pricing* (Rev. ed.). Princeton University Press.

Cooper, I., & Priestley, R. (2011). Real investment and risk dynamics. *Journal of Financial Economics*, *101*(1), 182–205. https://doi.org/10.1016/j.jfineco.2011.02.002

Daniel, K., Hirshleifer, D., & Sun, L. (2020). Short- and long-horizon behavioral factors. *Review of Financial Studies*, *33*(4), 1673–1736. https://doi.org/10.1093/rfs/hhz069

Euromedia Research Group. (2022). Euromedia ownership monitor: Sweden [Report].

Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, *33*(1), 3–56. https://doi.org/10.1016/0304-405X(93)90023-5

Fama, E. F., & MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, *81*(3), 607–636. https://doi.org/10.1086/260061

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics*, *127*(2), 645–700. https://doi.org/10.1093/qje/qjs001

Gourio, F. (2008). Disasters and recoveries. *American Economic Review: Papers & Proceedings*, *98*(2), 68–73. https://doi.org/10.1257/aer.98.2.68

Hirshleifer, D., Mai, D., & Pukthuanthong, K. (2025a). War discourse and disaster premium: 160 years of evidence from the stock market. *Review of Financial Studies*, *38*(2), 457–506. https://doi.org/10.1093/rfs/hhae081

Hirshleifer, D., Mai, D., & Pukthuanthong, K. (2025b). War discourse and the cross section of expected stock returns. *Journal of Finance*, *80*(6), 3589–3638. https://doi.org/10.1111/jofi.13482

Jacobs, H., & Müller, S. (2020). Anomalies across the globe: Once public, no longer existent? *Journal of Financial Economics*, *135*(1), 213–230. https://doi.org/10.1016/j.jfineco.2019.06.004

Jensen, T. I., Kelly, B., & Pedersen, L. H. (2023). Is there a replication crisis in finance? *Journal of Finance*, *78*(5), 2465–2518. https://doi.org/10.1111/jofi.13249

Krishnamurthy, A., & Vissing-Jorgensen, A. (2012). The aggregate demand for Treasury debt. *Journal of Political Economy*, *120*(2), 233–267. https://doi.org/10.1086/666526

Lehmann, B. N., & Modest, D. M. (1988). The empirical foundations of the arbitrage pricing theory. *Journal of Financial Economics*, *21*(2), 213–254. https://doi.org/10.1016/0304-405X(88)90061-X

Lewellen, J., Nagel, S., & Shanken, J. (2010). A skeptical appraisal of asset pricing tests. *Journal of Financial Economics*, *96*(2), 175–194. https://doi.org/10.1016/j.jfineco.2009.09.001

Manela, A., & Moreira, A. (2017). News implied volatility and disaster concerns. *Journal of Financial Economics*, *123*(1), 137–162. https://doi.org/10.1016/j.jfineco.2016.01.032

Mehra, R., & Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, *15*(2), 145–161. https://doi.org/10.1016/0304-3932(85)90061-3

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, *55*(3), 703–708. https://doi.org/10.2307/1913610

Rietz, T. A. (1988). The equity risk premium: A solution. *Journal of Monetary Economics*, *22*(1), 117–131. https://doi.org/10.1016/0304-3932(88)90172-9

Shanken, J. (1992). On the estimation of beta-pricing models. *Review of Financial Studies*, *5*(1), 1–33. https://doi.org/10.1093/rfs/5.1.1

Shleifer, A., & Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, *52*(1), 35–55. https://doi.org/10.1111/j.1540-6261.1997.tb03807.x

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323. https://doi.org/10.1007/BF00122574

Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance*, *68*(3), 987–1035. https://doi.org/10.1111/jofi.12018

Appendix A. Summary of Variables

Table 7 summarizes the primary data sources used in the analysis, including sample periods, and key descriptive statistics for each variable.

Table 7. Summary of Variables

Variable	Freq.	Start	End	Min	Max	Mean	Std
War News Share	Wkly	1993	2024	0.019	0.344	0.046	0.024
Stock Returns	Wkly	1993	2024	−0.795	2.039	0.002	0.080
Risk-Free Rate, 1M T-Bill	Wkly	1993	2024	−0.000	0.002	0.000	0.001
Market Index, SIX Return Index	Wkly	1993	2024	−0.205	0.186	0.003	0.028
Government Bond Yield, 2Y	Wkly	1993	2024	−0.879	10.408	2.795	2.746
Government Bond Yield, 5Y	Wkly	1993	2024	−0.704	11.470	3.205	2.762
IG Corporate Bond Index	Daily	2016	2024	217.580	249.540	228.436	6.790
NVIX War Component	Mthly	1993	2016	−0.254	0.721	0.058	0.108
Geopolitical Risk Index	Mthly	1993	2024	39.046	512.530	100.612	47.878
FF4 Factors (SHoF)	Wkly	1993	2019	—	—	—	—

This table summarizes the primary data sources. All statistics are computed over the 1993–2024 sample period (subject to data availability). War News Share is the fraction of total articles classified as war-related across six major Swedish outlets (Retriever media archive). Stock Returns cover Swedish listed stocks, reported as weekly returns prior to liquidity filters. The Risk-Free Rate is the Swedish one-month T-bill rate from Riksbanken. The Market Index is the SIX Return Index. Government bond yields are Riksbanken constant-maturity par yields. The IG Corporate Bond Index is the S&P Sweden Investment Grade Corporate Bond Index, available from 2016. The NVIX War Component is from Manela and Moreira (2017), available through 2016. The Geopolitical Risk Index is from Caldara and Iacoviello (2022). FF4 Factors are the published Fama–French–Carhart series from the Swedish House of Finance (Aytug et al., 2020), available through December 2019.

Appendix B. Boolean Keyword String

War-related articles are identified in the Retriever media archive using the following Boolean keyword string, applied simultaneously across all six monitored outlets. The string is case-insensitive. The four included terms exhaust the core Swedish vocabulary for armed conflict, spanning colloquial (*krig* [war], *invasion* [invasion]) and formal-register expressions (*militär konflikt* [military conflict], *väpnad konflikt* [armed conflict]) to maximize recall across the relevant semantic field. The exclusion list follows the false-positive filtering logic of Caldara and Iacoviello (2022) and addresses a more systematic problem specific to Swedish morphology, in which *krig* and *invasion* are productive compounding roots that generate numerous high-frequency non-military terms, the omission of which materially improves the signal-to-noise ratio of the war share.

Included terms:

```
("krig*" OR "invasion*" OR "militär konflikt*" OR "väpnad konflikt*")
```

Excluded terms (ANDNOT logic):

```
ANDNOT (priskrig* OR handelskrig* OR kulturkrig* OR valutakrig* OR budkrig* OR
budgivningskrig* OR reklamkrig* OR tullkrig* OR lönekrig* OR hyreskrig* OR
bostadskrig* OR rekryteringskrig* OR talangkrig* OR techkrig* OR mediekrig* OR
streamingkrig* OR varumärkeskrig* OR influencerkrig* OR "PR-krig*" OR
prispresskrig* OR lågpriskrig* OR rabattkrig* OR konkurrenskrig* OR
marknadskrig* OR "kriget om kunderna*" OR "kriget om marknadsandelar*" OR
"kriget om talanger*" OR "kriget om jobben*" OR "kriget om makten*" OR "kriget
mot inflationen*" OR "kriget mot brottslighet*" OR "kriget mot cancer*" OR
"kriget mot klimatförändringar*" OR "kriget mot covid*" OR "kriget mot gäng*"
OR "krig mellan bolag*" OR "krig i branschen*" OR turistinvasion* OR
besökarinvasion* OR folkinvasion* OR "invasion av turister*" OR "invasion av
fans*" OR "invasion av människor*" OR "invasion av produkter*" OR "invasion av
marknaden*" OR "invasion av appar*" OR "invasion av migranter*" OR
migrantinvasion* OR invasionsart* OR artinvasion* OR "invasiva arter*" OR
"biologisk invasion*" OR "digital invasion*" OR "AI-invasion*" OR "kulturell
invasion*" OR "medial invasion*" OR "Netflix-invasion*" OR butiksinvasion*)
```

Appendix C. WarFac Construction Details

The rolling AR(1) persistence parameter $\hat{\rho}_t$ is bounded to $(-0.999, +0.999)$ throughout the estimation. If the rolling OLS estimate satisfies $|\hat{\rho}_t| \geq 0.999$, it is clipped to $\operatorname{sgn}(\hat{\rho}_t) \times 0.999$ and the intercept is recomputed as

$$\hat{\alpha}_t = \bar{S}^{\text{war}}_{[t-W,t-1]} - \hat{\rho}_t \bar{S}^{\text{war}}_{[t-W-1,t-2]}, \tag{14}$$

where bars denote window means, to preserve the in-sample mean.

Without this bound, a rolling window dominated by a major war-news spike can produce near-unit-root estimates of $\hat{\rho}_t$. A near-unit-root AR(1) then predicts persistently elevated war coverage in subsequent weeks; as coverage naturally fades after the spike, the actual series falls below these inflated predictions, generating large negative innovations that would be classified as de-escalation shocks despite reflecting return-to-mean from an outlier rather than genuine reductions in war-risk perception. The bound ensures that predicted values decay rather than persist near-explosively, so that all large negative *WarFac* innovations reflect genuine reductions in war-related media attention rather than mechanical artefacts of the rolling estimation window.

Appendix D. War Share and *WarFac* Diagnostics

Table 8 reports time-series diagnostics for both the Swedish war share and the *WarFac* innovation series. Table 9 reports distributional statistics and Table 10 reports pairwise Pearson correlations for *WarFac* alongside the Carhart (1997) four factors.

Table 8. War Share and *WarFac* Time-Series Diagnostics

Test	Statistic	p-value	Conclusion
Panel A: War share			
Augmented Dickey–Fuller	-6.978	< 0.001	Reject unit root
Ljung–Box	—	< 0.001	Strong autocorrelation
Rolling mean CV	0.277	—	Time-varying
Rolling variance CV	2.491	—	Time-varying
Panel B: WarFac innovations			
Ljung–Box	—	< 0.001	Reject white noise
AC(1)	-0.039	—	Near-zero
AC(2)	-0.012	—	Near-zero
AC(3)	$+0.044$	—	Near-zero
AC(4)	$+0.075$	—	Near-zero
AC(5)	$+0.060$	—	Near-zero

Panel A reports time-series diagnostic tests for the Swedish war share over January 1990 to January 2024. The Augmented Dickey–Fuller test uses four lags. The Ljung–Box statistic reports the minimum p-value across ten lags. The rolling coefficient of variation (CV) measures parameter variation over a 156-week window. Panel B reports the Ljung–Box statistic and individual autocorrelation coefficients at lags 1 through 5 for the *WarFac* innovation series over January 1993 to January 2024.

Table 9. Distributional Statistics: *WarFac* and Carhart Factors

Factor	Mean (ann. %)	Std (ann. %)	Skewness	Ex. Kurtosis
WarFac	-0.98	6.23	4.421	60.468
MKT	11.00	20.61	-0.615	5.503
SMB	-2.78	12.57	-0.221	2.503
HML	5.16	14.57	-0.182	7.212
UMD	19.60	29.05	-0.194	2.989

This table reports distributional statistics for *WarFac* and the Carhart (1997) four factors over the estimation sample July 1993 to January 2024. Mean and standard deviation are annualized by multiplying weekly figures by 52 and $\sqrt{52}$ respectively. Excess kurtosis is reported relative to the normal distribution benchmark of zero.

Table 10. Pairwise Pearson Correlations: *WarFac* and Carhart Factors

	WarFac	MKT	SMB	HML	UMD
WarFac	1.000	0.022	−0.061	−0.040	−0.033
MKT	0.022	1.000	−0.357	−0.164	−0.174
SMB	−0.061	−0.357	1.000	−0.059	0.075
HML	−0.040	−0.164	−0.059	1.000	−0.268
UMD	−0.033	−0.174	0.075	−0.268	1.000

This table reports pairwise Pearson correlations between *WarFac* and the Carhart (1997) four factors over the estimation sample July 1993 to January 2024.

Appendix E. Factor Construction

This appendix describes the construction of the Fama–French–Carhart four factors at weekly frequency from Finbas data. Aytug et al. (2020) construct Swedish factors up until 2019, which is published by the Swedish House of Finance. We follow a similar approach, constructing Swedish Fama–French–Carhart four factors up until 2024.

Stock universe

The universe comprises all common equity listings on all Swedish stock markets with non-missing bid prices in the Finbas database.[10] Preferred shares, warrants, convertibles, and depository receipts are excluded. The market capitalization used for size breakpoints and value-weighting is the lagged market capitalization from the prior week. For companies reporting in foreign currencies, stock prices, market capitalizations and book equity are all converted to SEK in each week, using the average exchange rate for each year. Where a firm has multiple share classes listed simultaneously, returns are aggregated across share classes using a value-weighted approach.

Market factor (MKT)

The market excess return is the weekly return on the SIX Return Index minus the Swedish one-month T-bill rate from Riksbanken. The SIX Return Index is spliced from the Finbas series (1980–2019) and the Refinitiv Eikon series (2009–2026), with the Finbas series taking precedence over the overlap period (2009–2019).

Size factor (SMB)

At the beginning of each week, stocks are ranked by market capitalization and split into a small (bottom 80%) and big (top 20%) group using the 80th percentile as the breakpoint. SMB is the value-weighted return on the small group minus the value-weighted return on the big group, rebalanced weekly.

Value factor (HML)

At the beginning of each week, stocks with non-missing book equity are ranked by book-to-market ratio (book equity divided by market capitalization). Stocks are assigned to a low (bottom 30%), neutral (middle 40%), or high (top 30%) book-to-market group using the 30th and 70th percentiles of the ranked distribution as breakpoints. HML is the value-weighted average return of the high book-to-market stocks minus the value-weighted average return of the low book-to-market stocks, rebalanced weekly.

---

[10] Four administrative and non-trading categories are excluded throughout, identified by the Finbas codes `SSEUTL`, `SSEA2U`, `EXTERN`, and `INOFF` in line with Aytug et al. (2020).

Momentum factor (UMD)

At the beginning of each week, stocks are ranked by their cumulative return over weeks $t-52$ to $t-5$, skipping the most recent four weeks to avoid the short-term reversal effect. Stocks are assigned to a winner (top 10%) or loser (bottom 10%) group using the 10th and 90th percentiles as breakpoints. UMD is the value-weighted return of the winner group minus the value-weighted return of the loser group, rebalanced weekly.

Validation

All weekly stock returns are winsorised cross-sectionally at the 1st and 99th percentiles in each week prior to factor construction. The replicated factors are validated against the published SHoF weekly factor series of Aytug et al. (2020) over 1,136 weeks of common overlap from January 1993 to December 2019. The Pearson correlations between the replicated and published series are 0.998 (MKT), 0.852 (SMB), 0.621 (HML), and 0.725 (UMD).

The lower HML correlation likely reflects sensitivity to book-equity timing conventions, where minor differences in fiscal year-end matching and the treatment of missing book values can compound across the sample. As HML serves as a control factor rather than the factor of interest, this discrepancy should not materially affect the main results.

Appendix F. Full Coefficient Vector: Main Specification

Table 11 reports the complete estimated coefficient vector from the main specification of the Cochrane (2005) two-pass cross-sectional regression.

Table 11. Full Coefficient Vector: Individual Stock Main Specification (WarFac+FF4)

Parameter	$\hat{\lambda}$ (weekly)	Ann. (%)	t_{OLS}	t_{Shanken}
Zero-beta rate λ_0	+0.00106***	+5.51	+2.92***	+2.89***
λ_{WarFac}	−0.00049**	−2.53	−2.02**	−2.00**
λ_{MKT}	+0.00176***	+9.15	+4.19***	+4.15***
λ_{SMB}	−0.00154***	−8.01	−4.74***	−4.70***
λ_{HML}	−0.00028	−1.46	−0.78	−0.77
λ_{UMD}	+0.00301***	+15.65	+3.06***	+3.03***
Shanken c	1.020			
R^2	4.12%			
MAPE	0.270% per week			
N stocks	1,011			
Mean T	587 weeks			

This table reports the full estimated coefficient vector from the Cochrane (2005) full-sample two-pass cross-sectional regression of the main specification (WarFac+FF4) applied to 1,011 individual Swedish stocks over July 1993 to January 2024. The second-pass OLS regresses time-series mean excess returns on full-sample factor betas, including an intercept λ_0 following Black (1972). Ann. (%) annualizes the weekly coefficient by multiplying by 52. t_{OLS} is the unadjusted OLS t-statistic; t_{Shanken} applies the Shanken (1992) errors-in-variables correction with Shanken factor $c = 1.020$. MAPE is the mean absolute pricing error from the second-pass regression. Mean T is the average number of weekly return observations per stock used in the first pass. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Appendix G. Robustness of the *WarFac* Risk Premium

Table 12 reports robustness checks on the *WarFac* risk premium across three panels. Panel A varies the *WarFac* construction. Panel B varies the minimum observation threshold for the individual stock universe. Panel C reports the external benchmark validation. [11]

Table 12. Robustness of the *WarFac* Risk Premium

Variant	Ann. (%)	t_{Shanken}
Panel A: WarFac construction		
Baseline (rolling AR(1), 3-year window)	−2.53**	−2.00
Rolling 2-year window	−2.89**	−2.28
Rolling 5-year window	−2.14*	−1.70
Rolling ARMA(1,1), 3-year window	−7.21***	−3.24
Panel B1: Cross-sectional $\hat{\lambda}$ (Cochrane two-pass, individual stocks)		
Baseline (EW, 260-week min, all caps)	−2.53**	−2.00
Min obs = 3 years (156 weeks)	−4.45***	−3.99
Panel B2: FMP_{stock} (D1−D10), FF4 α		
Baseline (EW, 260-week min, all caps)	+2.09	0.42
Min obs = 3 years (156 weeks)	+6.48*	1.71
Value-weighted deciles	+0.81	0.10
Panel C: External benchmark validation		

Benchmark	Pearson r	N (months)
GPR (Caldara & Iacoviello, 2022)	0.588^{***}	494
GPR Threats (GPR_T)	0.684^{***}	494
GPR Acts (GPR_A)	0.424^{***}	494
NVIX War (Manela & Moreira, 2017)	0.513^{***}	374

Panels A and B1 report the annualized *WarFac* risk premium and Shanken (1992)-corrected t-statistic from the Cochrane (2005) two-pass regression (WarFac+FF4) applied to individual stocks. Panel B2 reports the annualized FF4 alpha and Newey–West t-statistic (26 lags) for $\text{FMP}_{\text{stock}} = \text{D1}-\text{D10}$. Panel C reports Pearson correlations of the monthly Swedish war share against the GPR index and its subindices (Caldara & Iacoviello, 2022) and the NVIX War component (Manela & Moreira, 2017), with peak lead-lag correlation at lag zero for both benchmarks. The baseline uses a rolling AR(1) *WarFac* with a 156-week window, equal-weighted portfolios, and a 260-week minimum observation threshold. Panel A replaces the baseline *WarFac* series with three alternatives. Panel B varies the minimum observation threshold and portfolio weighting. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

[11] A rolling AR(2) specification was estimated but excluded from the robustness set. At $r = 0.680$ with the baseline, the AR(2) residual is materially different from the AR(1) innovation, reflecting the second lag stripping out pricing-relevant medium-frequency dynamics alongside serial correlation. The resulting premium is economically negligible and statistically insignificant, better interpreted as a diagnostic that over-filtering the war share removes the signal that investors price.

Table 13 reports the Cochrane two-pass result applied to the 92-portfolio cross-section as test assets.

Table 13. Portfolio Robustness: Cochrane Two-Pass on 92 Double-Sorted Sector Portfolios

Parameter	$\hat{\lambda}$ (weekly)	t_{OLS}	t_{Shanken}	Shanken c	R^2
$\hat{\lambda}_{WarFac}$	−0.00277*	−2.08	−1.91*	1.185	33.77%

This table reports the Cochrane (2005, ch. 12) full-sample two-pass cross-sectional regression result for the main specification (WarFac+FF4) applied to 92 double-sorted sector portfolios over July 1993 to January 2024. The 92 portfolios are formed by combining 11 equal-weighted GICS sector portfolios with three sets of 27 portfolios constructed by double-sorting each sector against size, book-to-market, and momentum terciles respectively, subject to a minimum of five stocks per cell. t_{Shanken} applies the Shanken (1992) errors-in-variables correction. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Appendix H. Identified Shock Events

Table 14 lists all 18 positive shock events identified at the 2σ threshold after applying the 13-week de-clustering window, ordered chronologically. Each event is characterized by its date and the corresponding geopolitical episode.

Table 14. Identified Positive Shock Events (2σ Threshold)

Date	Geopolitical episode
May 1995	Gulf War aftermath and Srebrenica
March 1999	NATO Kosovo campaign
December 1999	Second Chechen War escalation
June 2001	Pre-September 11 escalations
October 2001	September 11 and aftermath
March 2003	Iraq War
July 2006	Lebanon War
August 2008	Russia–Georgia war
May 2011	Arab Spring
March 2014	Crimea annexation
October 2014	Related escalations
November 2015	Paris attacks
December 2016	ISIS escalations
January 2020	Iran–USA tensions
August 2021	Afghanistan withdrawal
January 2022	Ukraine-related escalations
December 2022	Renewed Russia–Ukraine escalations
January 2024	Middle East conflict escalations

This table lists the 18 positive shock events identified at the 2σ threshold after applying a 13-week de-clustering window. Where multiple candidate events fall within 13 weeks of each other, only the week with the largest $|WarFac_t|$ is retained. One negative event (April 1994) is identified but excluded from the CAR analysis as a single observation does not support bootstrap inference.

Appendix I. Bond Market Robustness

Table 15 reports the sensitivity of the bond market *WarFac* coefficients to the choice of Newey–West HAC lag truncation parameter.

Table 15. Bond Market Response to *WarFac*: Newey–West Lag Length Sensitivity

	$\hat{\beta}_{WarFac}$ (t_{NW} in parentheses)		
Asset	NW = 4 lags	NW = 13 lags	NW = 26 lags
Panel A: Government bonds			
2Y Government Bond	−0.00842 (−1.04)	−0.00842 (−1.06)	−0.00842 (−1.08)
5Y Government Bond	−0.04305 (−1.99)**	−0.04305 (−2.06)**	−0.04305 (−2.09)**
Panel B: Term structure			
5Y−2Y Spread	−0.03463 (−2.42)**	−0.03463 (−2.54)**	−0.03463 (−2.58)***
Panel C: Corporate bonds			
IG Corporate Bond	−0.01317 (−0.81)	−0.01317 (−0.67)	−0.01317 (−0.60)

This table reports the sensitivity of the bond market *WarFac* coefficients to the choice of Newey–West HAC lag truncation parameter. Each cell reports $\hat{\beta}_{WarFac}$ from a time-series OLS regression of weekly asset excess returns on *WarFac* and the market excess return (MKT), with an intercept included, re-estimated at three lag choices: 4, 13, and 26 lags. The β point estimate is identical across lag choices within each asset by construction; only the standard error and resulting t-statistic vary. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Appendix J. AI-appendix

Tools Used and How

Claude and ChatGPT were used in a supportive capacity throughout the project. AI assisted with structuring and optimizing code for efficiency and readability, and was used for proofreading and improving the conciseness of written text.

Contribution to Thesis Quality

For the empirical pipeline, AI helped accelerate debugging and enforce consistency across the pipeline. For the written text, it helped identify grammatical errors.

Risks and Mitigating Measures

AI suggestions can introduce subtle errors or phrasings that pass initial inspection. We mitigated this by returning to reference papers as ground truth for methodological decisions and validating empirical outcomes, as well as by relying on supervisor oversight throughout.

Insights from the Process

AI proved most useful for well-specified technical tasks and less reliable for open-ended judgments. Analytical and interpretive decisions should be made independently to ensure best outcomes.