Geopolitical Risk and U.S. Investment-at-Risk*

Fabio Anobile† Francesco Frangiamore‡ Marco Maria Matarrese§
Jamel Saadaoui¶

Abstract

Geopolitical risk predicts a deterioration in U.S. investment-at-risk. A 30-point increase in historical GPR is associated with a 1.8-percentage-point decline in the fifth conditional percentile of next-quarter investment growth, compared with 0.15 percentage points at the mean. Median and upper-tail estimates are close to zero. The pattern survives financial and macroeconomic controls. Monthly time-varying local projections identify responses under the restriction that other economic shocks do not affect GPR within the month. Identified shocks persistently lower the fifth percentile relative to the median, with larger responses during high-GPR episodes.

Keywords: geopolitical risk; investment-at-risk; quantile regression; time-varying local projections; real investment

JEL codes: C22; E22; E32; F52

1 Introduction

Geopolitical tensions can delay investment even when their average macroeconomic effect appears modest. With costly reversibility, uncertainty raises the option value of waiting, so firms postpone projects until information improves (Dixit and Pindyck, 1994; Bloom, 2009). Firm-level evidence also links geopolitical risk (GPR) to lower capital expenditure, especially when capital is difficult to reverse (Wang et al., 2024). This mechanism implies a distributional prediction: GPR may worsen unusually adverse investment outcomes without shifting the center of the investment distribution by the same amount.

The investment literature links uncertainty to inaction when capital is partly irreversible (Bloom et al., 2007). Capital spending declines around elections (Julio and Yook, 2012), with policy uncertainty (Gülen and Ion, 2016), and with firm-specific political risk (Hassan et al., 2019). Caldara and Iacoviello (2022) document lower investment following geopolitical innovations, while Wang et al. (2024) emphasize capital irreversibility. We examine where the aggregate investment-growth distribution deteriorates beyond these mean and firm-heterogeneity results.

Investment-at-risk measures the severity of an adverse investment outcome at a fixed conditional probability. It can reveal worsening downside risk even when a mean forecast changes little. The fifth-percentile-minus-median gap further distinguishes a deteriorating lower tail from a common shift in the distribution. For capital-formation surveillance, a stable central forecast can therefore conceal a worsening adverse scenario.

*An earlier version circulated as SSRN Working Paper No. 5536200. We thank Leonardo Iania, Davide Furceri, Domenico Giannone, Tihana Skrinjaric, Francisco Serranito, and participants at the 27th INFER Annual Conference, the Bank of England seminar, and the INFER webinar, as well as the organizers and participants of the 1st Piraeus International Economic Conference, for helpful comments. Any errors are our own.

†LUM University, Bari, Italy. Email: anobile.phdstudent@lum.it.

‡University of Palermo, Palermo, Italy. Email: francesco.frangiamore@unipa.it.

§University of Foggia, Foggia, Italy. Email: marco.matarrese@unifg.it.

¶Corresponding author. Paris 8 University, France. Email: jamelsaadaoui@gmail.com.

Quarterly quantile regressions characterize predictive associations. Monthly TVP-LPs trace responses of fitted investment-risk measures to recursively identified GPR innovations. Their structural interpretation rests on a within-month exclusion motivated by Caldara and Iacoviello (2022), defended in Section 4 and Appendix A.5.

Other channels include war-related trade disruption (Glick and Taylor, 2010), trade-policy uncertainty (Handley and Limão, 2017), and political risk premia (Pástor and Veronesi, 2013). Threats may raise the option value of waiting, whereas acts combine uncertainty with physical, trade, and financial disruption. Component comparisons, financial controls, and durable-expenditure results can distinguish mechanism-consistent patterns, but no single comparison identifies a mechanism in isolation.

Adrian et al. (2019, 2022) show how financial conditions reshape downside growth risk across horizons. Loria et al. (2025) study fitted macroeconomic tail-risk responses to identified shocks. Furceri et al. (2025) document that several economic, financial, and uncertainty-related predictors have a stronger impact on upside risks to public debt. Reynolds and Lloyd (2026) estimate panel quantile LPs for 18 advanced and nine emerging economies over 1997Q4–2025Q1, examining GDP, inflation, trade, and yields. We isolate national-accounts investment, use a U.S. postwar sample beginning in 1947, compare threats with acts, and estimate time-varying investment-risk responses. The contribution concerns both the concentration of GPR's predictive information in adverse investment outcomes and the subsequent propagation of downside risk.

Higher GPR predicts lower investment-at-risk, with little association at the mean, median, or upper tail. The pattern survives financial and macroeconomic controls. Monthly TVP-LPs identify a persistent downside deterioration under the maintained restriction. Attenuation of the quarterly coefficient after 1971 motivates allowing propagation to vary over time.

2 Data and method

The underlying quarterly sample runs from 1947Q1 to 2025Q1. Real gross private domestic investment is the U.S. Bureau of Economic Analysis series distributed by FRED.[1] We use $y_t = 100\Delta \log(I_t)$. Real GPDI is the appropriate outcome because the paper concerns national-accounts capital formation and inventory investment; purchases of financial assets reallocate claims and are not investment in this sense. The historical GPR index (GPRH) of Caldara and Iacoviello (2022) starts in 1900. We convert the monthly index to quarterly frequency by averaging its three monthly observations. The U.S. provides both a long, internally consistent investment series and direct exposure to global geopolitical developments, but our evidence need not generalize to other economies.

Following the at-risk approach of Adrian et al. (2019), we estimate the one-quarter-ahead conditional quantile regression

$$Q_\tau(y_{t+1} \mid g_t, y_t) = \alpha_\tau + \beta_\tau g_t + \gamma_\tau y_t, \tag{1}$$

For $\tau = 0.05, 0.10, \ldots, 0.95$, g_t denotes GPR and β_τ measures its predictive association with the corresponding conditional quantile, holding current investment growth fixed. No causal interpretation is assigned to these quarterly coefficients. We report 90% confidence intervals using the quantile-regression HAC estimator of Galvao and Yoon (2024) and the smooth extended tapered block bootstrap of Gregory et al. (2018). Related work on testing growth vulnerabilities and macro-financial linkages also emphasizes time-series inference in quantile regression (Furno et al., 2025). We also simulate a VAR(4) in investment growth and GPR to obtain the coefficient bands implied by a linear data-generating process. Robustness specifications add the Baa–Aaa corporate credit spread and current GDP growth.

Quarterly GPDI provides a long sample spanning infrequent disruptions. Estimated monthly accounts permit structural responses of fitted conditional quantiles and interquantile gaps to be

[1] FRED series GPDIC1; billions of chained 2017 dollars, seasonally adjusted at an annual rate.

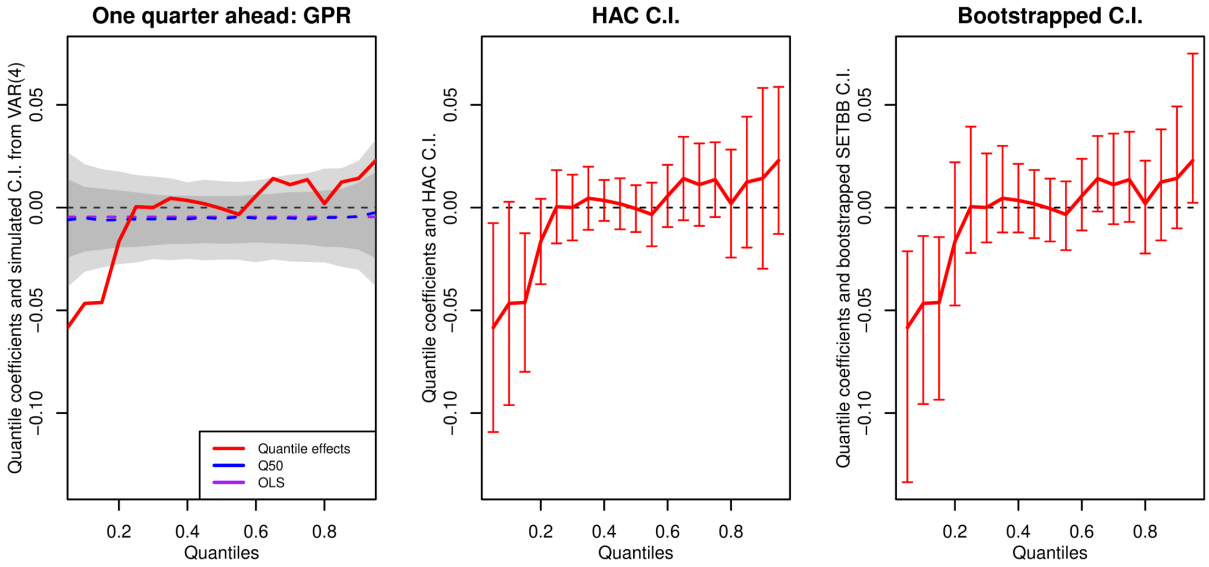


Figure 1: GPR and the conditional quantiles of next-quarter investment growth. The red line reports $\widehat{\beta}_\tau$ from equation (1). The left panel compares the quantile profile with bands simulated from a linear VAR(4); the middle and right panels report 90% HAC and block-bootstrap confidence intervals, respectively.

traced over monthly horizons. Identification is imposed in this second exercise.

3 Results

The linear benchmark gives a GPR coefficient of -0.005 with a HAC standard error of 0.009. At the fifth percentile, $\widehat{\beta}_{0.05}$ is approximately -0.06 (Figure 1). A common 30-point increase in GPRH thus predicts a 1.8-percentage-point lower fifth percentile of next-quarter growth, compared with 0.15 percentage points at the mean, holding current growth fixed. The fifth percentile measures severity at a fixed 5% conditional probability. Median and upper-tail estimates are close to zero; lower-tail coefficients fall outside the linear-model bands. HAC and bootstrap inference give the same qualitative ranking.

The asymmetry is not eliminated by observable financial or cyclical conditions. When the Baa–Aaa spread is included, GPR remains most negative in the lower tail; adding current GDP growth produces a similar profile (Online Appendix, Figures A.2–A.3). A specification using a long-run financial-stress proxy yields the same conclusion (Figures A.4–A.6). Distribution regressions provide an estimator-level check: higher GPR is associated with a greater probability of growth falling below severe negative thresholds, with no precise association at favorable thresholds (Figure A.20).

The long sample is informative but not innocuous. In the 1971–2025 subsample, the fifth-percentile estimate remains negative but is smaller and its 90% confidence intervals include zero, while financial conditions become more informative (Figures A.7–A.9). Removing the earliest Korean War observations also attenuates the coefficient (Figure A.10). This instability motivates the time-varying structural analysis below: an average full-sample coefficient can conceal large responses in high-tension periods.

Figure 2 separates geopolitical threats from realized acts. Threats have more negative lower-tail point estimates, although the formal comparison does not reject equality.[2] The point-estimate

[2] Appendix A.8 and Figure A.26 report the test of $H_0 : \beta_\tau^T - \beta_\tau^A = 0$. The comparison uses separate specifications on the same sample and 1,000 paired smooth extended tapered block-bootstrap replications, preserving the covariance between coefficient estimates. The 90% pointwise intervals include zero at every estimated quantile. The joint test at the fifth and tenth percentiles also fails to reject equality ($W = 0.390$, bootstrap $p = 0.843$; Appendix A.8). Non-rejection does not establish equality, and the decomposition remains suggestive rather than conclusive evidence for the proposed mechanism.

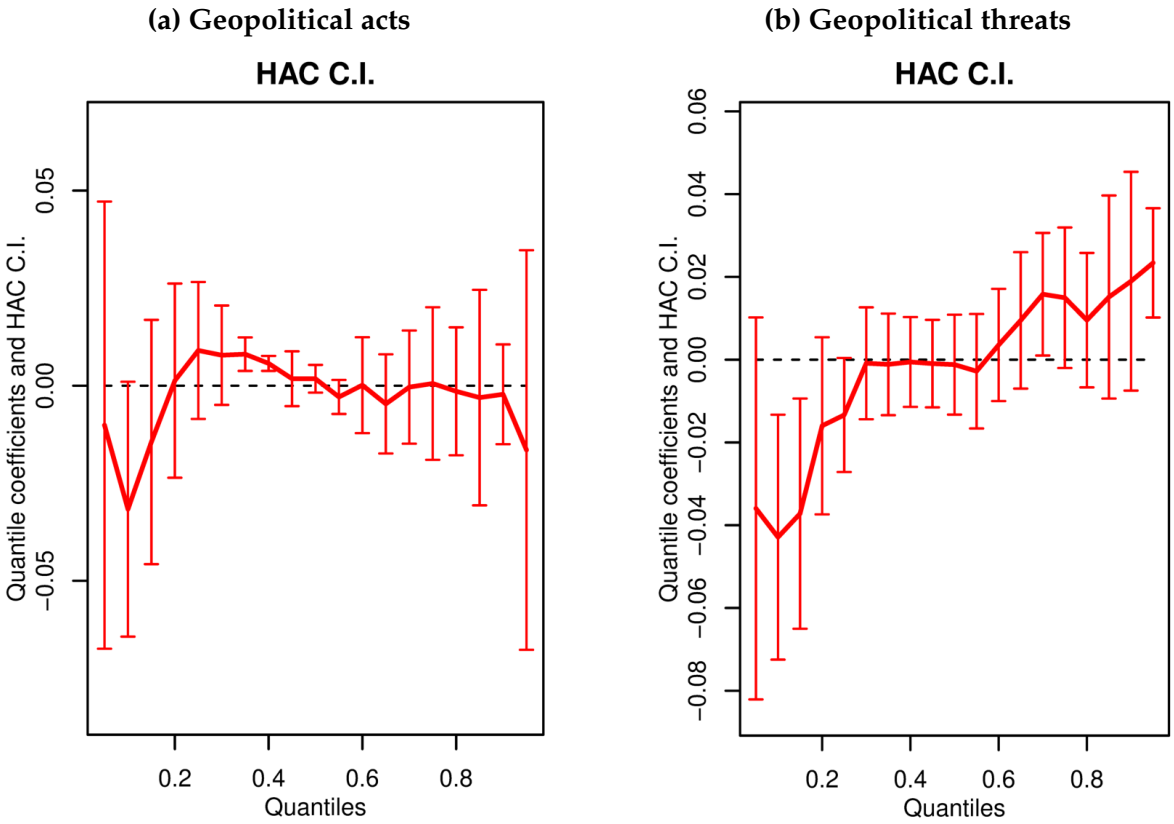


Figure 2: Quantile coefficients for GPR acts and threats. Panels report point estimates and 90% HAC confidence intervals from separate versions of equation (1).

ranking is consistent with a wait-and-see mechanism, but cannot distinguish it conclusively from other channels. Supplementary durable-expenditure results are likewise mechanism-consistent patterns rather than mechanism identification (Figures A.21–A.24).

4 Structural time-varying responses

Monthly dynamics use Barrett's (2025) estimated U.S. national accounts over 1990M1–2025M1. GPR is ordered first: geopolitical innovations may immediately affect investment, volatility, and sentiment, while other economic shocks do not contemporaneously affect GPR conditional on the included history. This applies Caldara and Iacoviello's (2022) quarterly ordering over a shorter interval. Sudden terrorist attacks and military escalations can immediately alter investment plans; monthly U.S. investment or sentiment disturbances are less plausible determinants of their timing. Feedback is permitted at longer horizons.

GPR's event-based construction distinguishes geopolitical coverage from adverse economic news. Caldara and Iacoviello (2022, pp. 1209–1212) validate its distinct content and examine news-crowding effects. Geopolitical crises may raise GPR and volatility together, consistently with our ordering. Their monthly Granger tests do not detect lagged economic predictors, complementing rather than testing the contemporaneous exclusion. Appendix A.5 formalizes and defends this restriction. Short-run identification applies to LPs as well as VARs (Plagborg-Møller and Wolf, 2021).

Following Loria et al. (2025), we first estimate one-month-ahead investment-growth quantiles and then trace how these fitted distributional objects respond to the identified shock. We use the path estimator of Inoue et al. (2024):

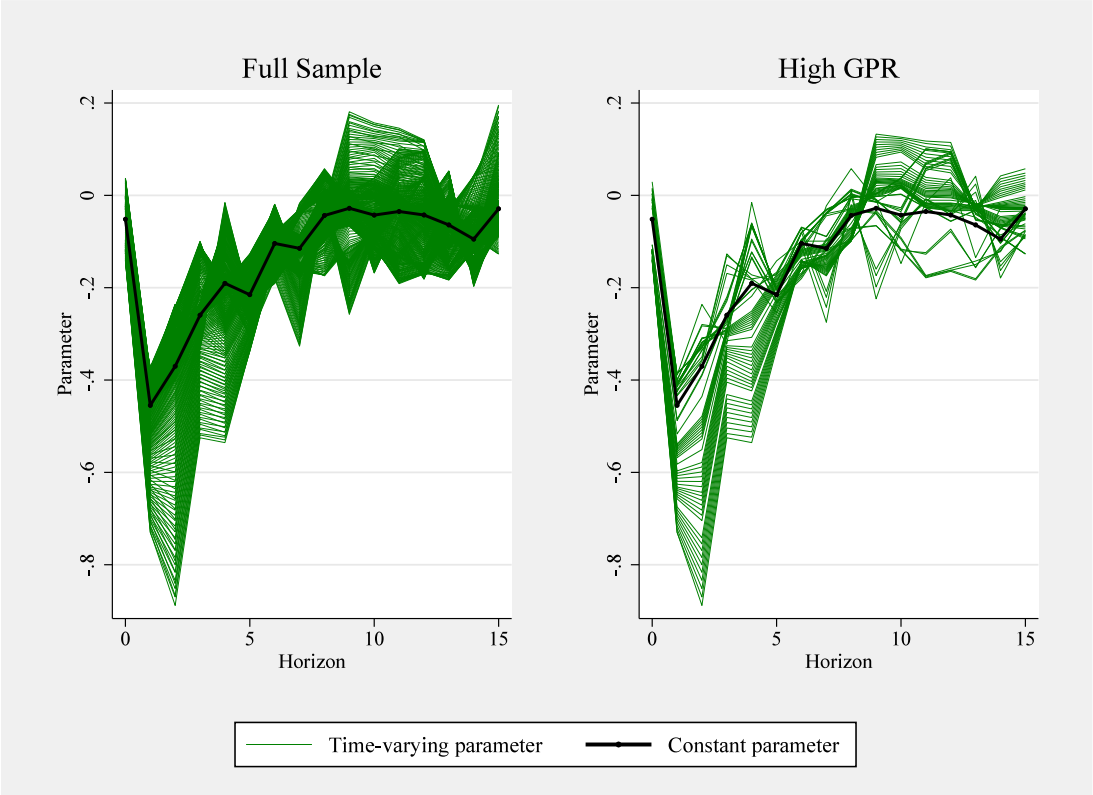


Figure 3: Time-varying responses of the fifth-percentile-minus-median investment-growth gap to a one-standard-deviation monthly GPR shock. Green lines are date-specific responses; the black line is the constant-parameter response. The right panel retains dates when GPR exceeds its 75th percentile. This selection is descriptive, not a separately estimated threshold model. Horizons are months.

$$q_{\tau,t+h} = \alpha_{\tau,h,t+h} + \beta_{\tau,h,t+h}\varepsilon_t^{gpr} + \sum_{j=1}^{24} \psi'_{\tau,h,j,t+h}\mathbf{w}_{t-j} + u_{\tau,t+h}, \tag{2}$$

Here $\varepsilon_t^{\text{gpr}}$ is normalized to a one-standard-deviation monthly GPR innovation. Controls comprise lags 1–24 of the fitted quantile, GPR, VIX, and Michigan sentiment. GPR lags condition on predictable history; the contemporaneous exclusion supplies identification. Current volatility and sentiment may transmit the shock. Coefficients vary across dates and horizons $h = 0, \ldots, 15$. Under the restriction, $\beta_{\tau,h,t+h}$ measures the structural response of fitted investment risk.[3]

A one-standard-deviation innovation lowers the fifth-percentile-minus-median gap by about 0.4 percentage points after one month in the constant-parameter specification; the response remains negative for roughly seven months (Figure 3). Selected high-GPR responses approach -0.9 percentage points at short horizons. This is an additional downside deterioration relative to the median. Tenth-percentile responses are weaker, and median-minus-95th-percentile responses much smaller (Figures A.12–A.19). The sharpest amplification occurs in the early 2000s; instability is clearest at short horizons. High-GPR paths describe historical variation, not a separately identified regime effect.

5 Conclusion

GPR contains information about adverse investment outcomes that is largely absent from central

[3]Following Loria et al. (2025), the first stage constructs a dated one-month-ahead conditional quantile; the second traces its response to the identified innovation. The fixed forecast horizon differs from the impulse-response horizon. The response is not a quantile treatment effect on realized growth.

forecasts. Monthly TVP-LPs, identified under a defended within-month exclusion, show a lower-tail deterioration relative to the median lasting several months. Larger dated responses coincide with high-tension episodes. Threats have more negative lower-tail point estimates than acts, but equality cannot be rejected, leaving the option-value-of-waiting interpretation suggestive. These findings complement Reynolds and Lloyd's (2026) cross-country tail-risk evidence by isolating capital formation as a distinct downside-risk margin. GPR can inform investment-risk surveillance even when mean forecasts move little.

Declarations

Competing interests. The authors declare no competing financial interests or personal relationships that could have influenced the work reported in this paper.

Funding. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data availability. The investment series is publicly available from FRED (GPDIC1), and the historical GPR index is publicly available from https://www.matteoiacoviello.com/gpr.htm. Series definitions and transformations are documented in the Online Appendix.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used OpenAI ChatGPT to assist with manuscript restructuring and language editing. After using this service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

Adrian, T., Boyarchenko, N., and Giannone, D. (2019). Vulnerable growth. *American Economic Review*, 109(4):1263–1289.

Adrian, T., Grinberg, F., Liang, N., Malik, S., and Yu, J. (2022). The term structure of growth-at-risk. *American Economic Journal: Macroeconomics*, 14(3):283–323.

Barrett, P. (2025). Estimated monthly national accounts for the United States. *IMF Working Papers*, 2025(134).

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.

Bloom, N., Bond, S., and van Reenen, J. (2007). Uncertainty and investment dynamics. *Review of Economic Studies*, 74(2):391–415.

Caldara, D. and Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4):1194–1225.

Dixit, A. K. and Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press, Princeton, NJ.

Furceri, D., Giannone, D., Kisat, F., Lam, W. R., and Li, H. (2025). Debt-at-risk. Centre for Economic Policy Research.

Furno, F., Giannone, D., Iania, L., Loria, F., and Schulz, C. (2025). Testing growth vulnerabilities and macro-financial linkages. Mimeo.

Galvao, A. F. and Yoon, J. (2024). HAC covariance matrix estimation in quantile regression. *Journal of the American Statistical Association*, 119(547):2305–2316.

Glick, R. and Taylor, A. M. (2010). Collateral damage: Trade disruption and the economic impact of war. *Review of Economics and Statistics*, 92(1):102–127.

Gregory, K. B., Lahiri, S. N., and Nordman, D. J. (2018). A smooth block bootstrap for quantile regression with time series. *Annals of Statistics*, 46(3):1138–1166.

Gülen, H. and Ion, M. (2016). Policy uncertainty and corporate investment. *Review of Financial Studies*, 29(3):523–564.

Handley, K. and Limão, N. (2017). Policy uncertainty, trade, and welfare: Theory and evidence for China and the United States. *American Economic Review*, 107(9):2731–2783.

Hassan, T. A., Hollander, S., van Lent, L., and Tahoun, A. (2019). Firm-level political risk: Measurement and effects. *Quarterly Journal of Economics*, 134(4):2135–2202.

Inoue, A., Rossi, B., and Wang, Y. (2024). Local projections in unstable environments. *Journal of Econometrics*, 244(2):105726.

Julio, B. and Yook, Y. (2012). Political uncertainty and corporate investment cycles. *Journal of Finance*, 67(1):45–83.

Loria, F., Matthes, C., and Zhang, D. (2025). Assessing macroeconomic tail risk. *Economic Journal*, 135(665):264–284.

Pástor, L. and Veronesi, P. (2013). Political uncertainty and risk premia. *Journal of Financial Economics*, 110(3):520–545.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local projections and VARs estimate the same impulse responses. *Econometrica*, 89(2):955–980.

Reynolds, J. and Lloyd, S. (2026). Tension across the tail: The cross-country effects of geopolitical risk on macroeconomic tail risk. *Economics Letters*, 264:112927.

Wang, X., Wu, Y., and Xu, W. (2024). Geopolitical risk and investment. *Journal of Money, Credit and Banking*, 56(8):2023–2059.

Online Appendix to Geopolitical Risk and U.S. Investment-at-Risk

Fabio Anobile Francesco Frangiamore Marco Maria Matarrese
Jamel Saadaoui

This appendix reports data definitions, robustness evidence, and the complete time-varying local-projection results referenced in the main text.

A.1 Data and transformations

The underlying quarterly sample is 1947Q1–2025Q1. Real gross private domestic investment is FRED series GPDIC1, sourced from the U.S. Bureau of Economic Analysis: billions of chained 2017 dollars, seasonally adjusted at an annual rate. Investment growth is $100[\log(I_t) - \log(I_{t-1})]$. The historical GPR index of Caldara and Iacoviello (2022) is monthly and begins in 1900; the quarterly observation is the arithmetic mean of the three monthly values. Real GDP is transformed analogously. The corporate credit spread is the quarterly average of Moody's seasoned Baa yield minus the corresponding Aaa yield. The National Financial Conditions Index is converted from weekly to quarterly frequency by averaging observations within each quarter.

A.2 Financial and macroeconomic controls

Figure A.2 estimates the baseline quantile regression after adding the Baa–Aaa corporate credit spread. The upper row reports the GPR coefficient; the lower row reports the spread coefficient. Figure A.3 additionally includes current real GDP growth. Across both specifications, the GPR profile remains most negative in the lower tail. These figures show that the predictive distributional association is not absorbed by the included observables; they do not establish causal identification for the quarterly regression.

A.3 Financial-stress estimates

No single financial-stress series spans the full quarterly sample. We therefore use the newspaper-based U.S. Financial Stress Indicator of Püttmann (2018) through 2016Q4 and the St. Louis Fed Financial Stress Index (FRED series STLFSI4) from 2017Q1 onward. The historical indicator is published with mean 100 and unit standard deviation; subtracting 100 places it on the zero-centered scale of STLFSI4. Weekly STLFSI4 observations are averaged within each quarter. Figure A.4 displays the component series and the resulting proxy.

Figure A.5 uses the financial-stress proxy in place of GPR. Figure A.6 adds the proxy to the baseline GPR specification. The GPR profile remains sharply negative in the lower tail, indicating that the investment-at-risk result is not absorbed by measured financial stress.

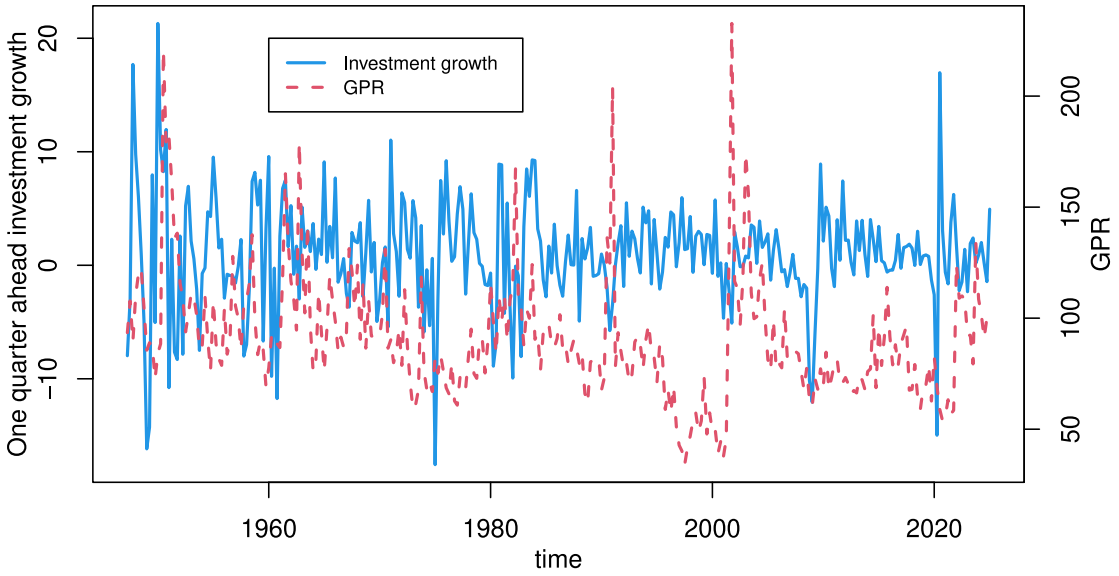


Figure A.1: Historical GPR and quarter-on-quarter real private investment growth. The blue solid line is investment growth (left axis); the red dashed line is GPR (right axis).

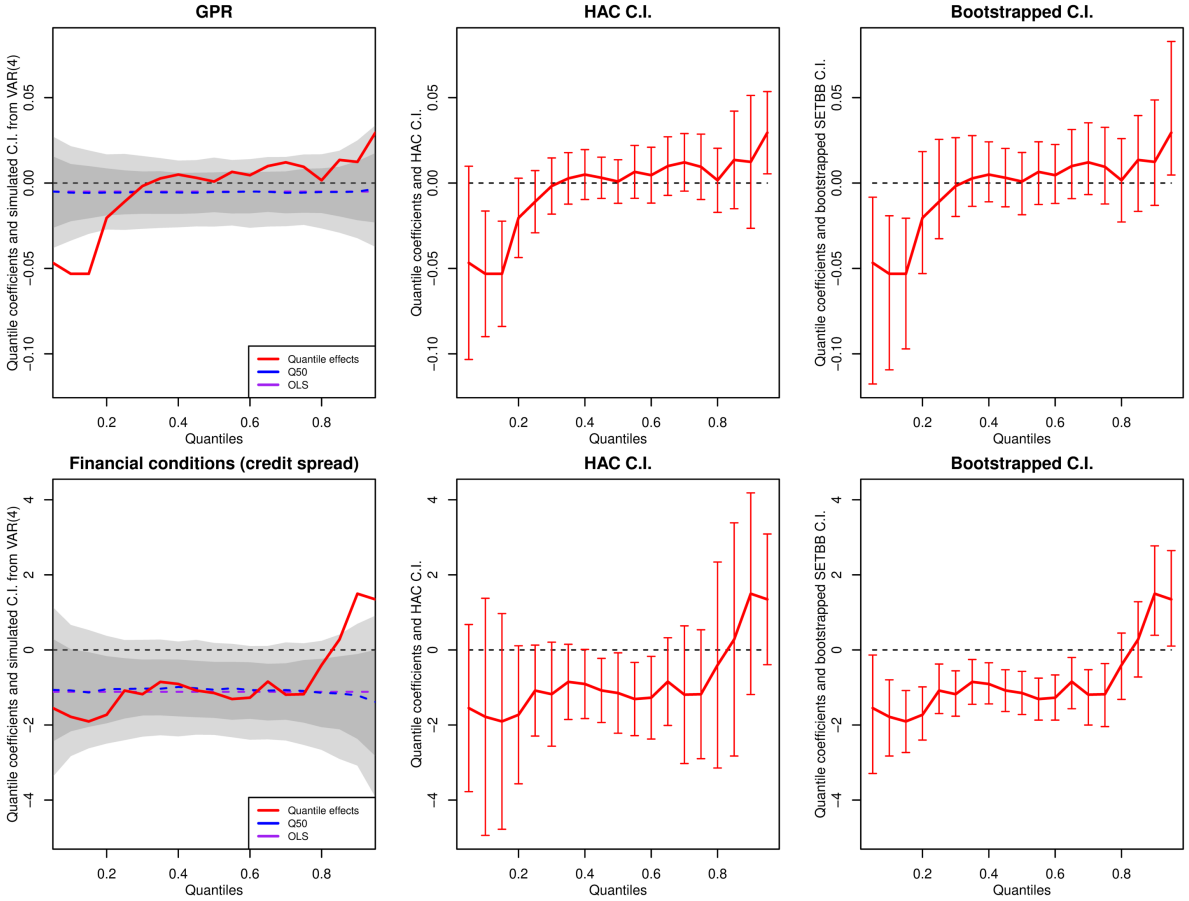


Figure A.2: Quantile coefficients for GPR with the credit spread included as a control.

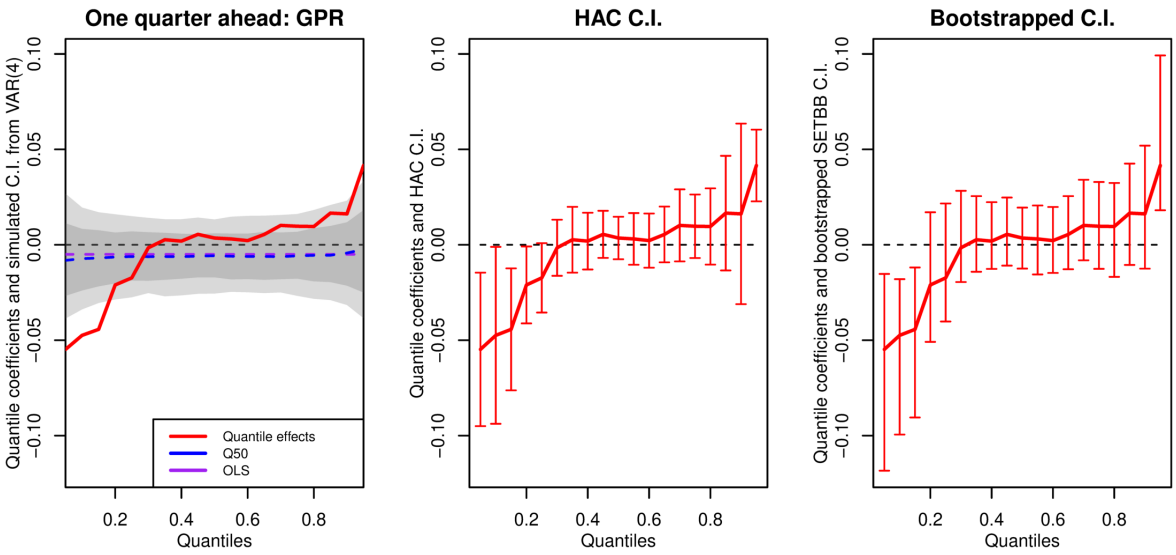


Figure A.3: GPR coefficients after controlling for the credit spread and current GDP growth.

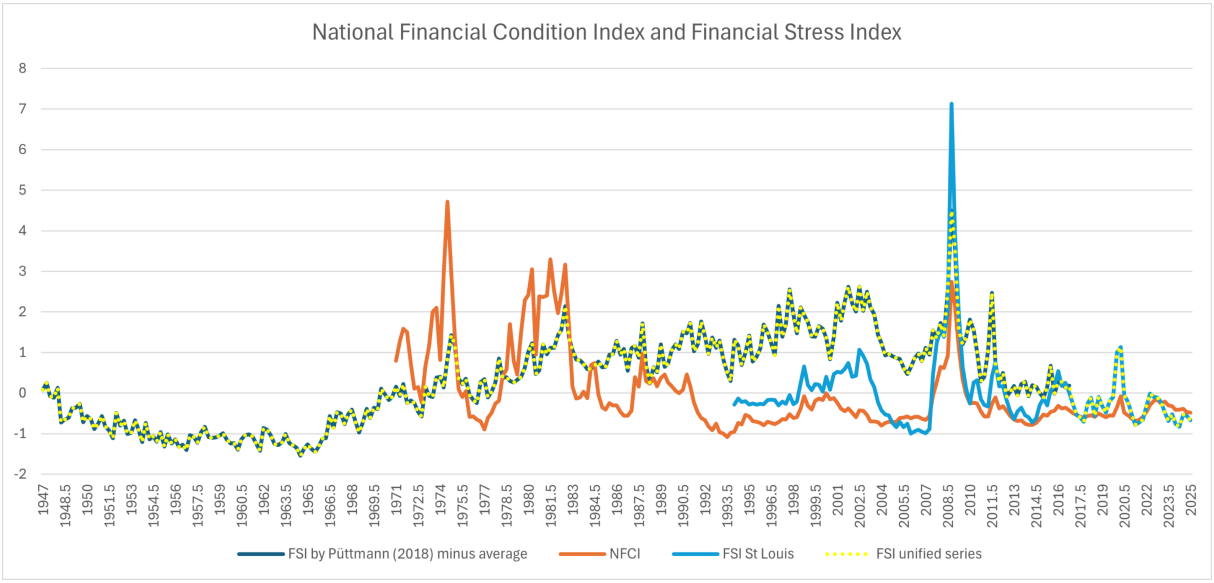


Figure A.4: Construction of the financial-stress proxy. The historical Püttmann indicator, centered by subtracting 100, is used through 2016Q4; the quarterly average of STLFSI4 is used from 2017Q1. The Chicago Fed NFCI is shown for comparison.

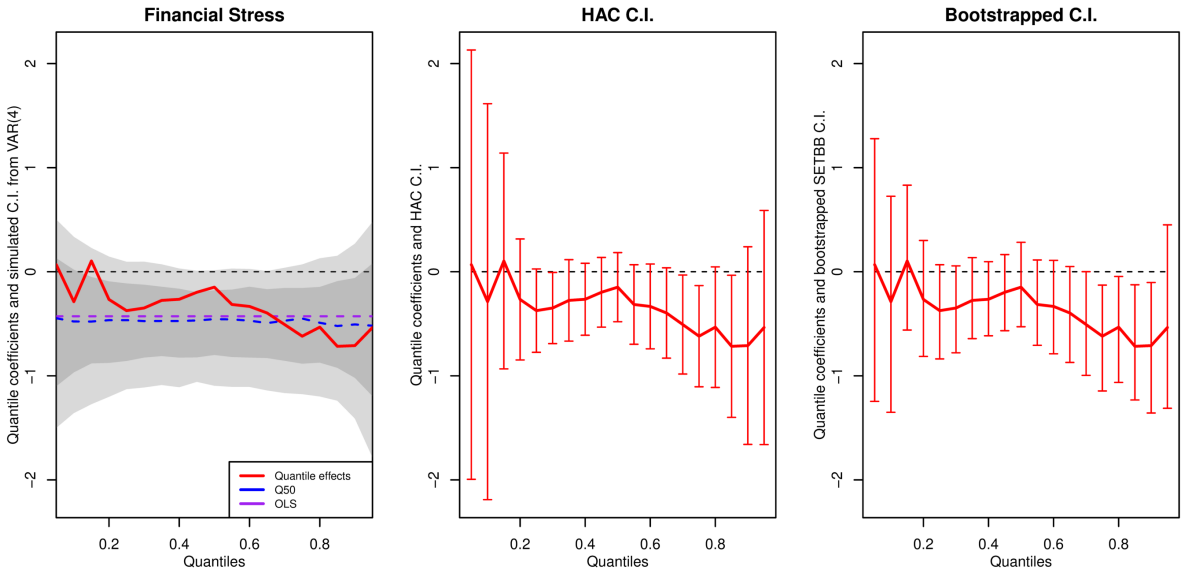


Figure A.5: Financial stress and the conditional quantiles of next-quarter investment growth.

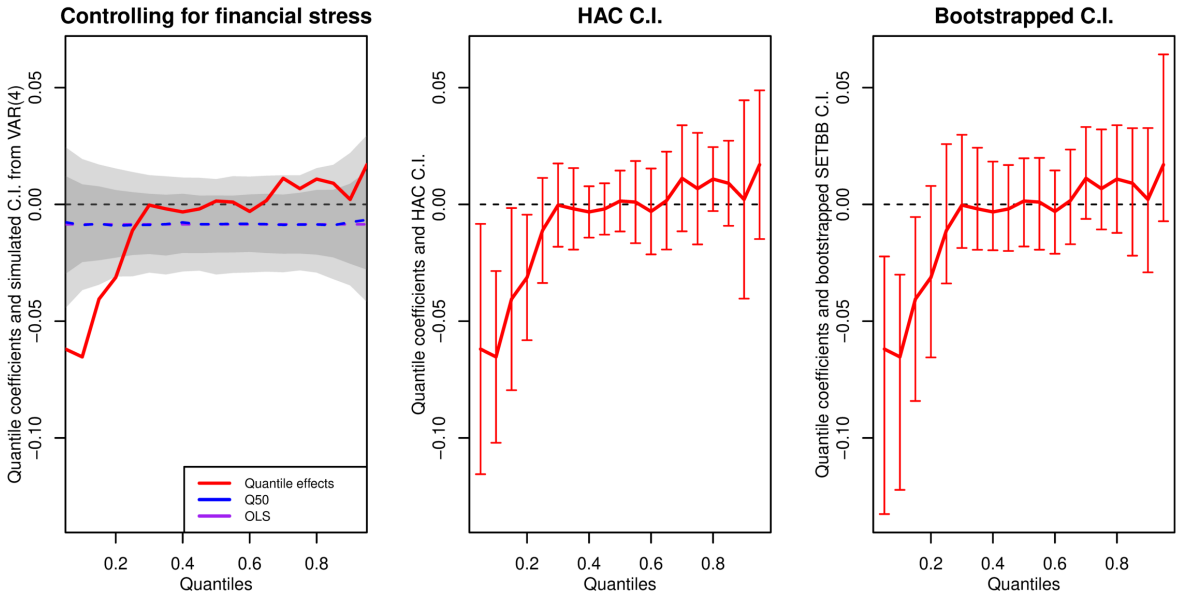


Figure A.6: GPR coefficients after controlling for the financial-stress proxy.

A.4 Sample sensitivity

The post-1971 evidence is an important qualification. Figure A.7 re-estimates the GPR specification over 1971–2025. The fifth-percentile coefficient remains negative but is less precisely estimated. Over the same period, the NFCI and the Baa–Aaa spread are informative about investment tails (Figures A.8–A.9). Figure A.10 drops the earliest Korean War observations by starting in 1952; Figure A.11 ends the sample before the COVID-19 pandemic. Together, these checks indicate that the full-sample GPR result is not driven by COVID-19 but is sensitive to the early postwar geopolitical episodes.

A.5 Structural time-varying local projections

The monthly exercise uses the estimated U.S. monthly national accounts of Barrett (2025). The sample is 1990M1–2025M1, beginning with the availability of the VIX. Monthly investment growth is the log difference multiplied by 100. In a first step, we estimate the conditional quantiles of one-month-ahead investment growth. The second step traces the response of those fitted quantiles and their interquantile gaps to GPR innovations, following the two-step quantile-local-projection design of Loria et al. (2025).

Structural identification imposes a monthly recursive exclusion with GPR ordered first. Let u_t^g denote the GPR innovation conditional on the included lagged information and u_t^x the other reduced-form innovations. The identifying representation is

$$\begin{pmatrix} u_t^g \\ u_t^x \end{pmatrix} = \begin{pmatrix} p_{gg} & 0' \\ p_{xg} & P_{xx} \end{pmatrix} \begin{pmatrix} \varepsilon_t^g \\ \varepsilon_t^x \end{pmatrix}, \qquad E(\varepsilon_t \varepsilon_t') = I. \qquad (A1)$$

With $p_{gg} > 0$, the normalized innovation $\varepsilon_t^g = u_t^g / p_{gg}$ identifies the geopolitical shock. Other variables may respond immediately through p_{xg}. The restriction excludes contemporaneous effects of other structural shocks on GPR; it does not require GPR and financial or real reduced-form innovations to be uncorrelated. It applies the direction of causation in Caldara and Iacoviello's (2022) quarterly U.S. BVAR over a shorter monthly interval and permits lagged feedback.

The exclusion is economically motivated by the timing of sudden geopolitical events. Terrorist attacks and military escalations can immediately change investment plans and financing conditions, whereas unexpected monthly movements in U.S. investment or sentiment are less plausible determinants of their timing. GPR's event-based search criteria and validation support its distinct geopolitical content (Caldara and Iacoviello, 2022, pp. 1209–1212). Their comparison with financial volatility includes geopolitical crises in which both indicators rise, consistent with the transmission permitted here. Their monthly Granger tests in Online Appendix B.5 concern lagged predictability and provide complementary evidence, not a test of contemporaneous exclusion.

For $\tau \in \{0.05, 0.10, 0.50, 0.95\}$ and horizons $h = 0, \ldots, 15$, we estimate

$$q_{\tau,t+h} = \alpha_{\tau,h,t+h} + \beta_{\tau,h,t+h} \varepsilon_t^{gpr} + \sum_{j=1}^{24} \psi_{\tau,h,j,t+h}' \mathbf{w}_{t-j} + u_{\tau,t+h}, \qquad (A2)$$

The GPR innovation is normalized to one standard deviation, and w_t collects the relevant fitted quantile, GPR, the VIX, and Michigan consumer sentiment. Lags 1 through 24 condition on the observed history of investment risk, geopolitical developments, volatility, and confidence. The GPR lags absorb the component of persistent coverage predictable from that history. In a conventional linear LP, this conditioning has a Frisch–Waugh–Lovell interpretation: the current-GPR coefficient uses variation remaining after the controls have been partialled out. This explains the role of lag conditioning, while the contemporaneous exclusion supplies the

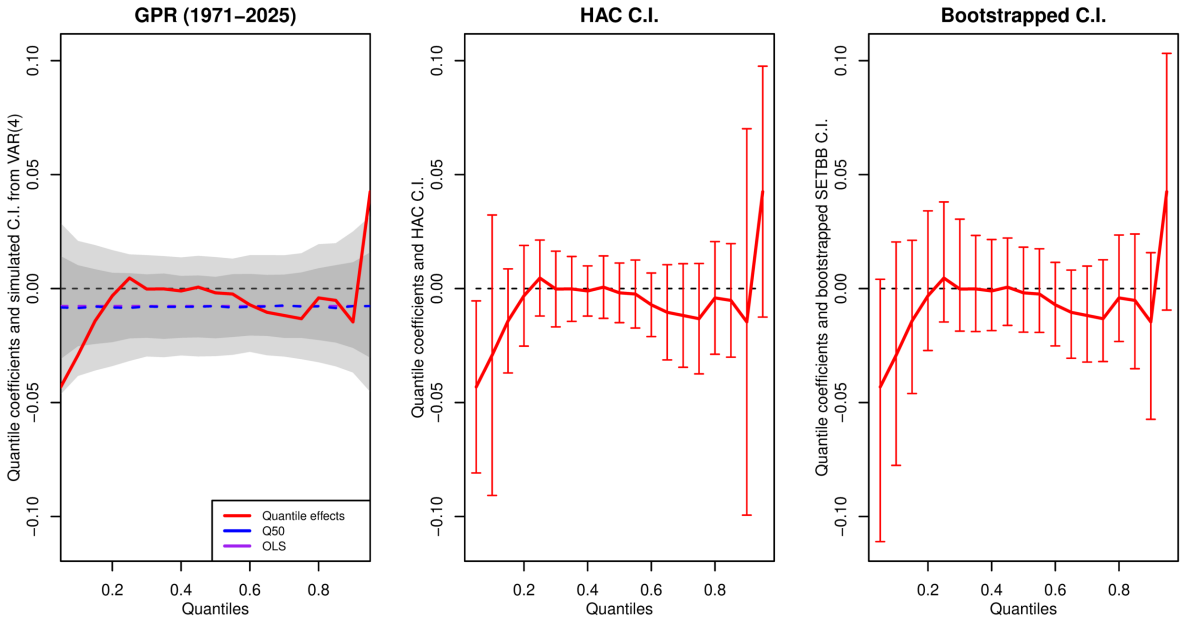


Figure A.7: GPR coefficients in the 1971–2025 subsample.

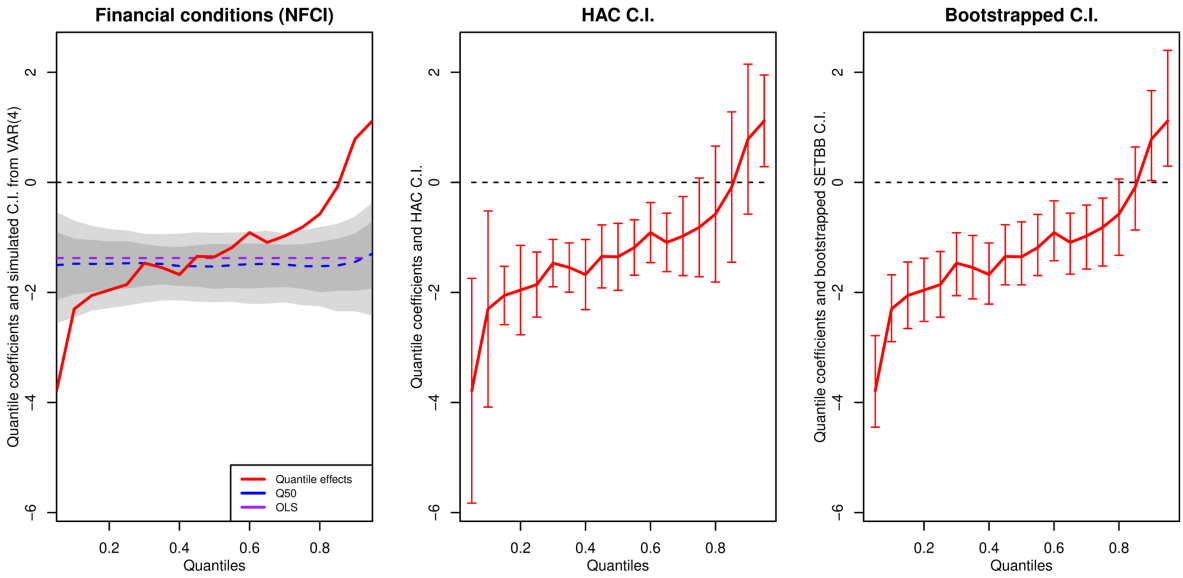


Figure A.8: NFCI coefficients in the 1971–2025 subsample.

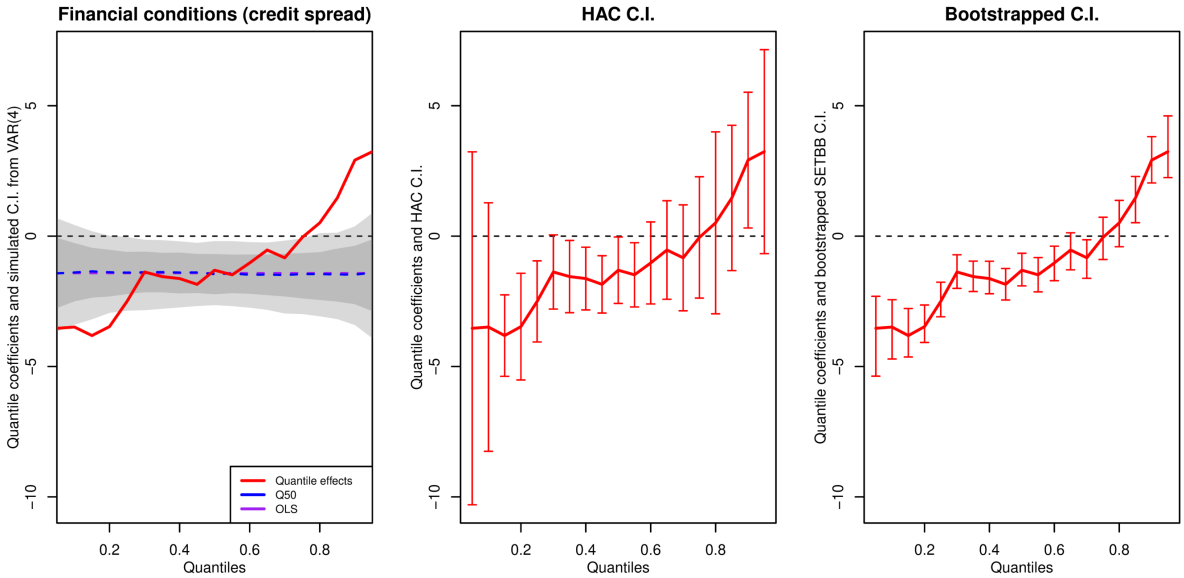


Figure A.9: Credit-spread coefficients in the 1971–2025 subsample.

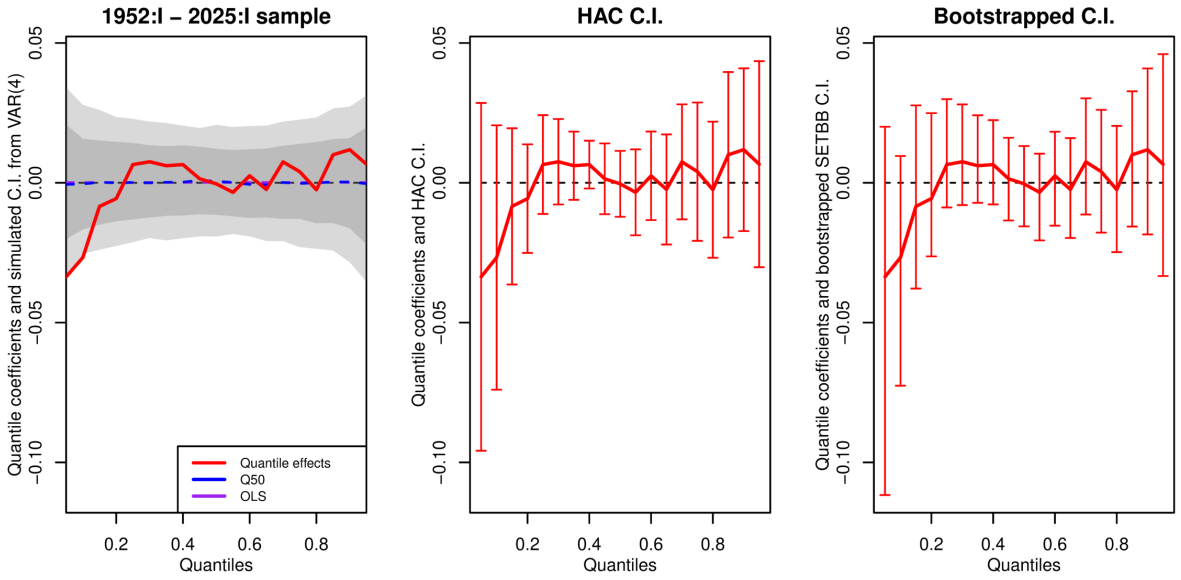


Figure A.10: Baseline GPR specification for the sample beginning in 1952.

structural interpretation. The time-varying path estimator of Inoue et al. (2024) governs how propagation changes across dates and horizons. Applying short-run restrictions to LPs is justified by Plagborg-Møller and Wolf (2021); it does not require a finite-sample identity between the path estimator and a VAR.

Current volatility and sentiment may transmit the geopolitical shock. Their lagged values control preceding conditions without conditioning away those contemporaneous responses. An omitted investment determinant is a confounder when it is also correlated with the identified GPR innovation after conditioning on the included history. A non-geopolitical disturbance entering both investment risk and measured GPR within the month would challenge the exclusion. For the sudden geopolitical events motivating the restriction, this contemporaneous confounding channel is less evident than the transmission from geopolitical developments to economic conditions. The restriction remains maintained and is supported by the timing argument and index-validation evidence.

The response concerns a fitted one-month-ahead conditional risk measure. The forecast horizon used to construct that measure remains fixed as the impulse-response horizon varies. It is therefore distinct from a quantile treatment effect on realized growth. We allow slopes and error variances to vary over calendar time. The time-path panels report horizons zero through five with 90% confidence bands. In the horizon-profile figures, the right panel descriptively selects paths observed when GPR exceeds its sample 75th percentile; it is not a separate threshold-model estimate.

Figures A.12–A.13 report the fifth-percentile response, and Figures A.14–A.15 the tenth-percentile response. Both decline over the first several months, with the fifth percentile moving more. The fifth-percentile-minus-median gap reaches about -0.4 percentage points after one month in the constant-parameter estimate and remains negative for roughly seven months; selected high-GPR paths approach -0.9 percentage points at short horizons (Figures A.16–A.17). The median-minus-95th-percentile gap is much smaller (Figures A.18–A.19). The largest deterioration occurs in the early 2000s. Because several 90% bands include the constant response, instability is strongest at short horizons rather than uniform across the entire response.

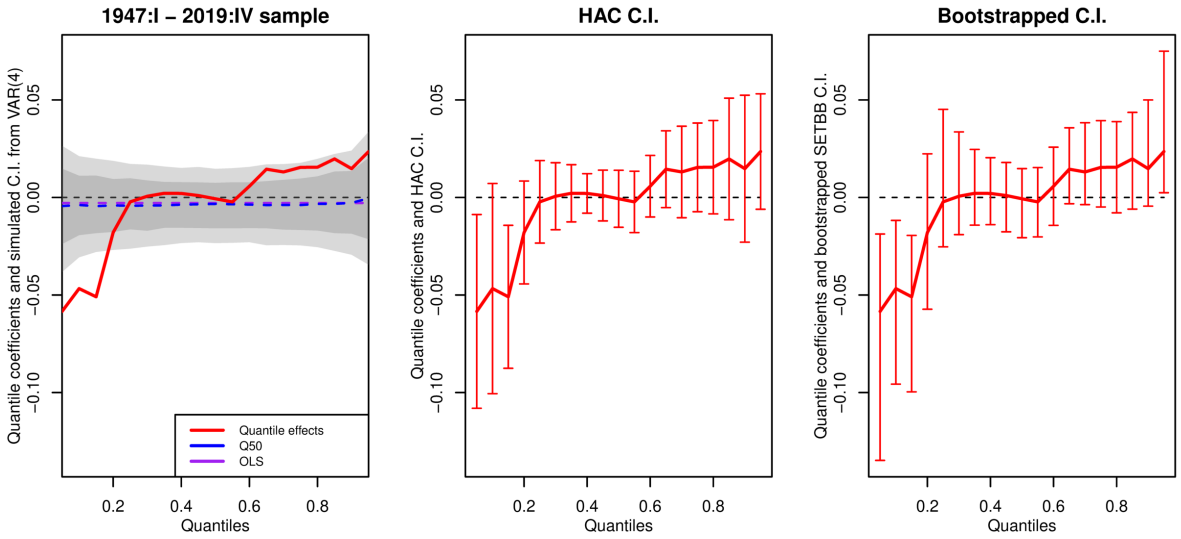


Figure A.11: Baseline GPR specification ending in 2019Q4.

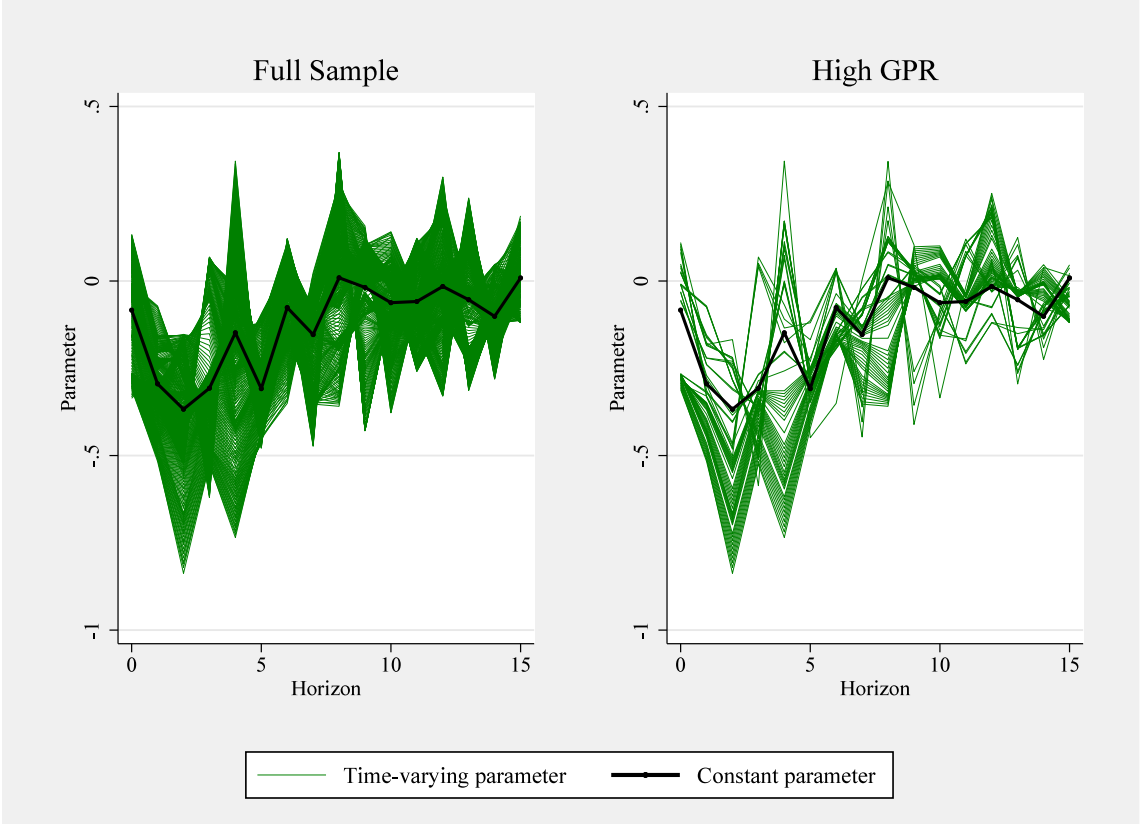


Figure A.12: Time-varying local projections for the fifth conditional percentile. Green lines are date-specific responses and the black line is the constant-parameter response. The right panel selects months in which GPR exceeds its 75th percentile.

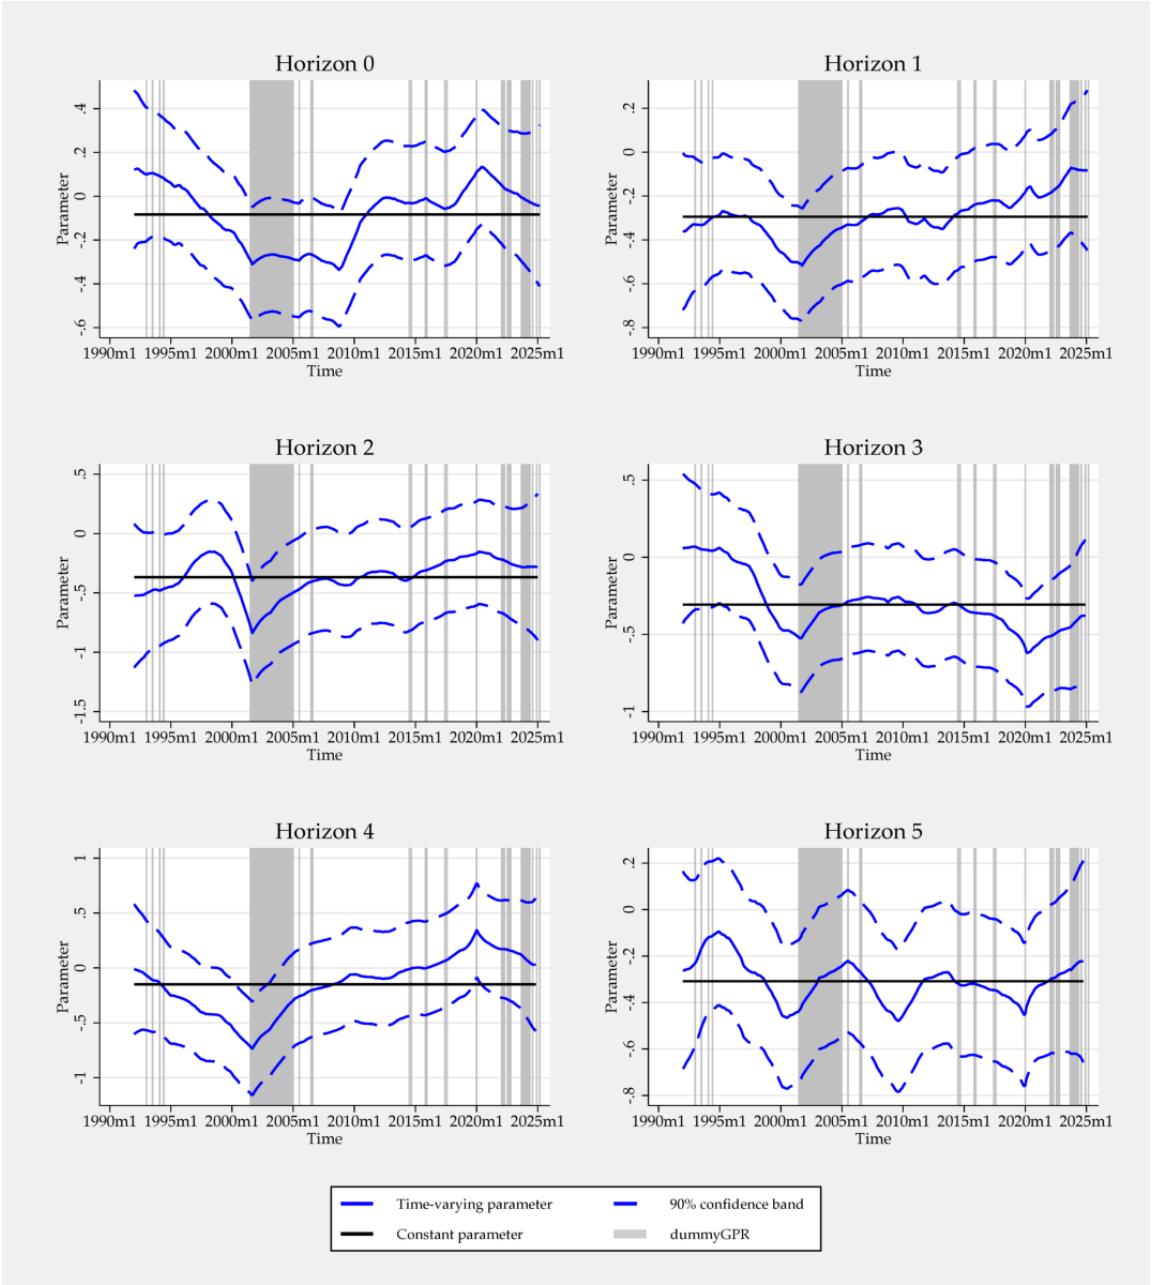


Figure A.13: Fifth-percentile response across time at horizons zero through five. Dashed lines are 90% confidence bands; gray bars identify months in which GPR exceeds its 75th percentile.

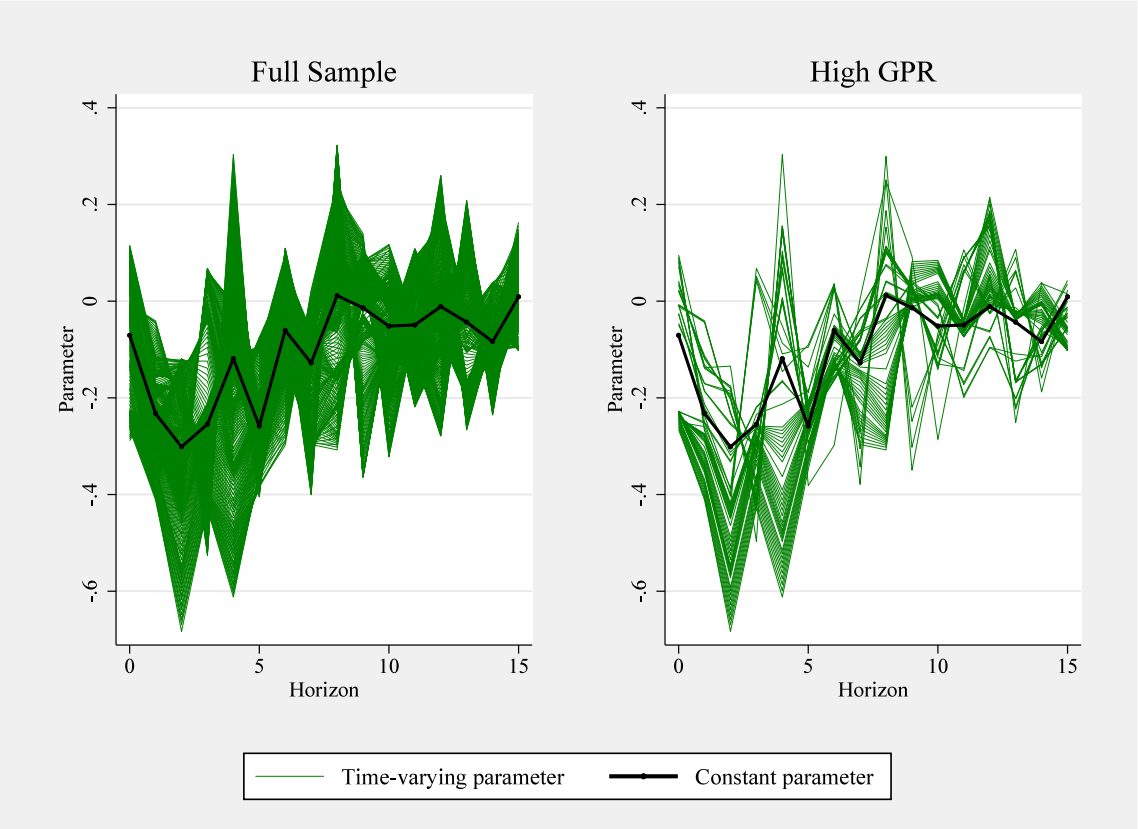


Figure A.14: Time-varying local projections for the tenth conditional percentile.

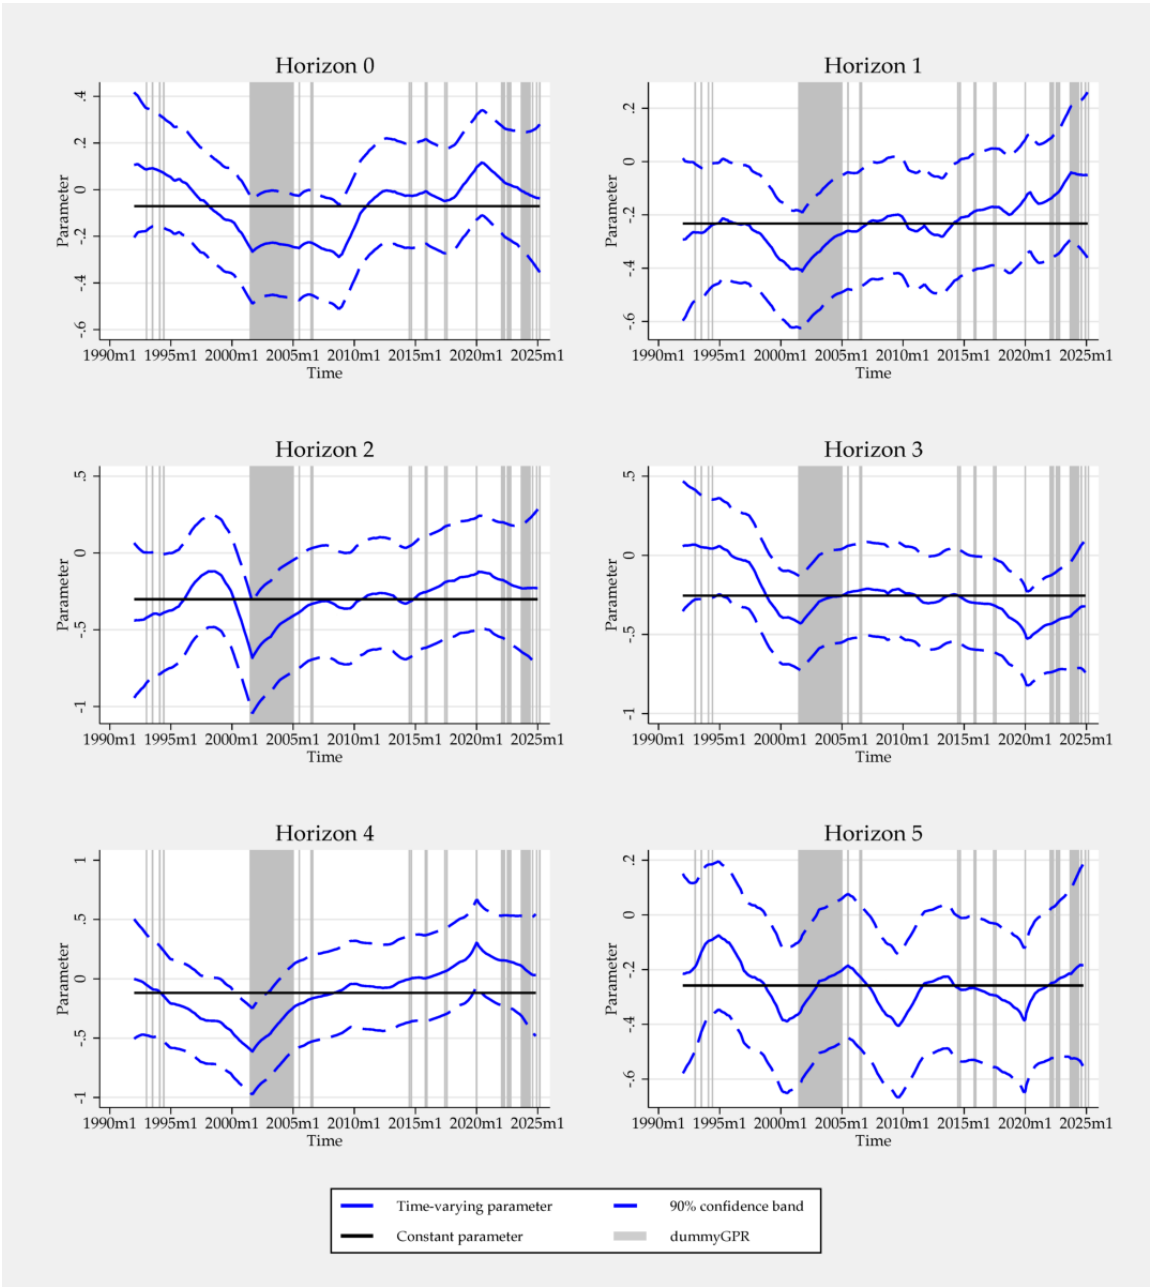


Figure A.15: Tenth-percentile response across time at horizons zero through five. Dashed lines are 90% confidence bands.

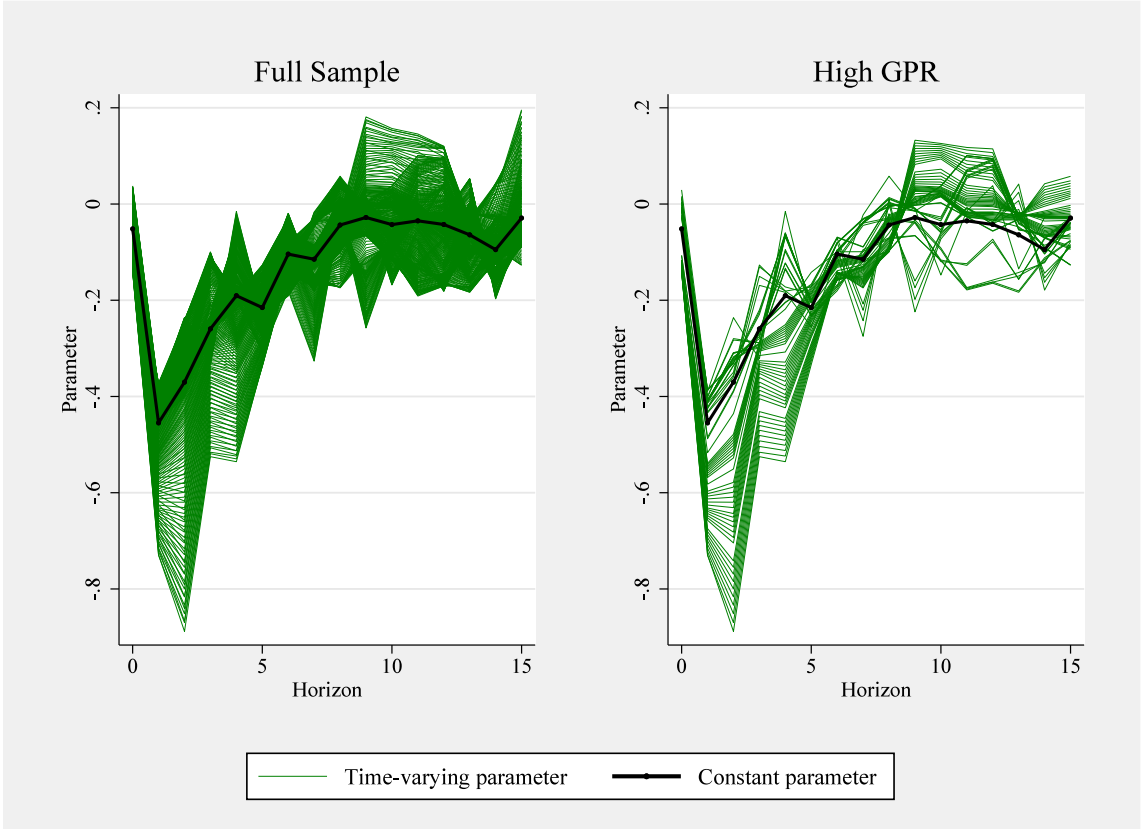


Figure A.16: Time-varying responses of the fifth-percentile-minus-median gap.

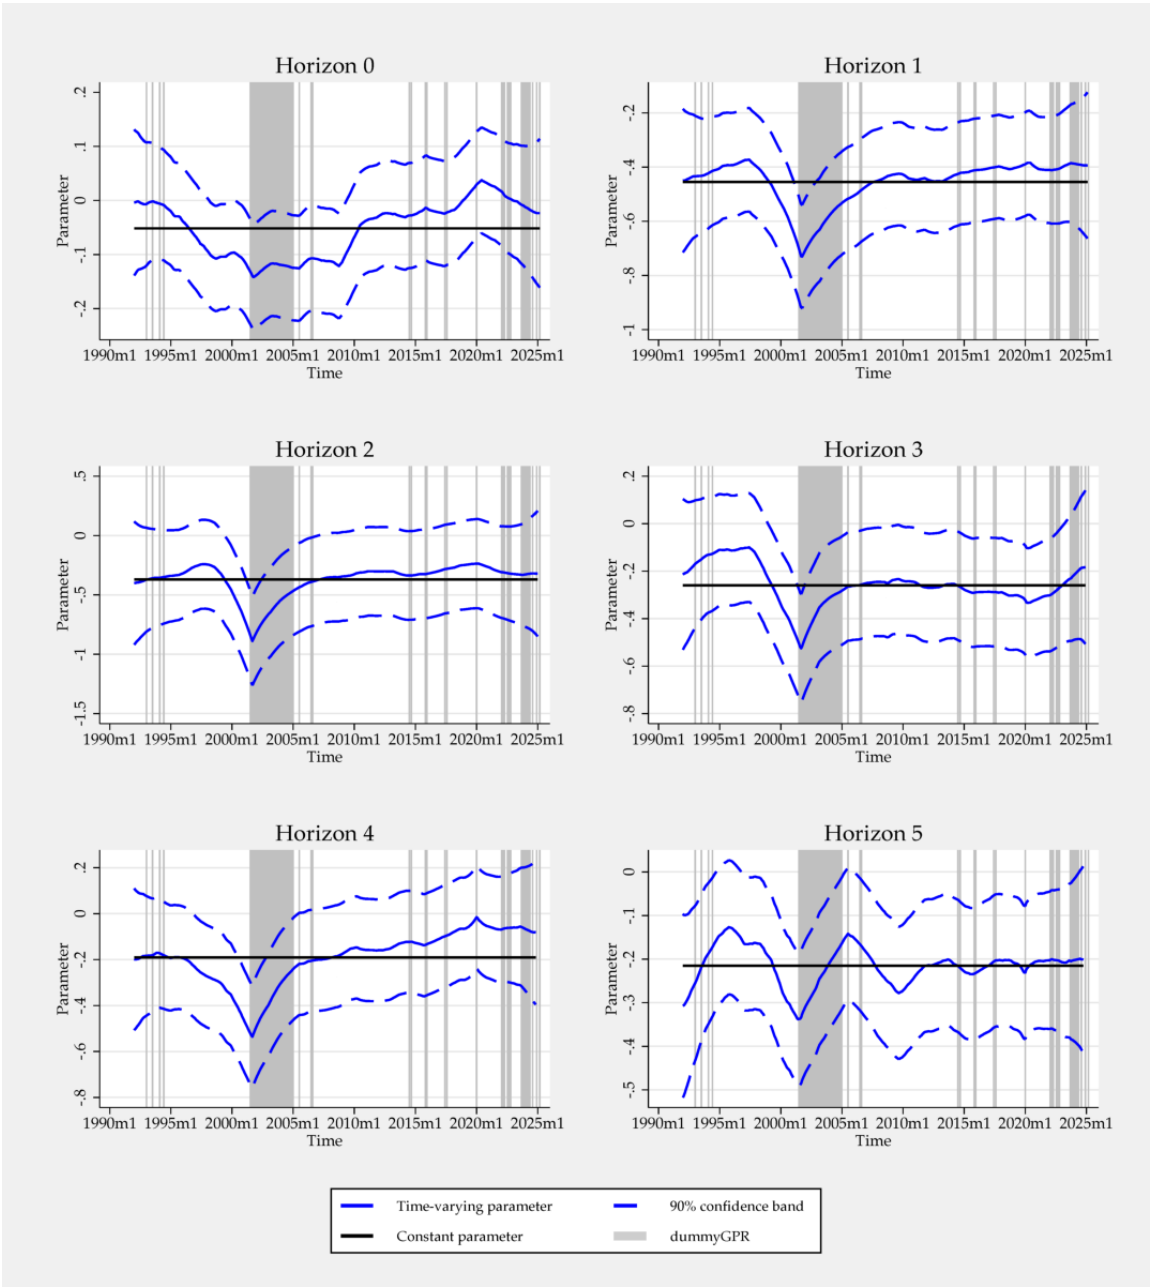


Figure A.17: Fifth-percentile-minus-median response across time at horizons zero through five. Dashed lines are 90% confidence bands.

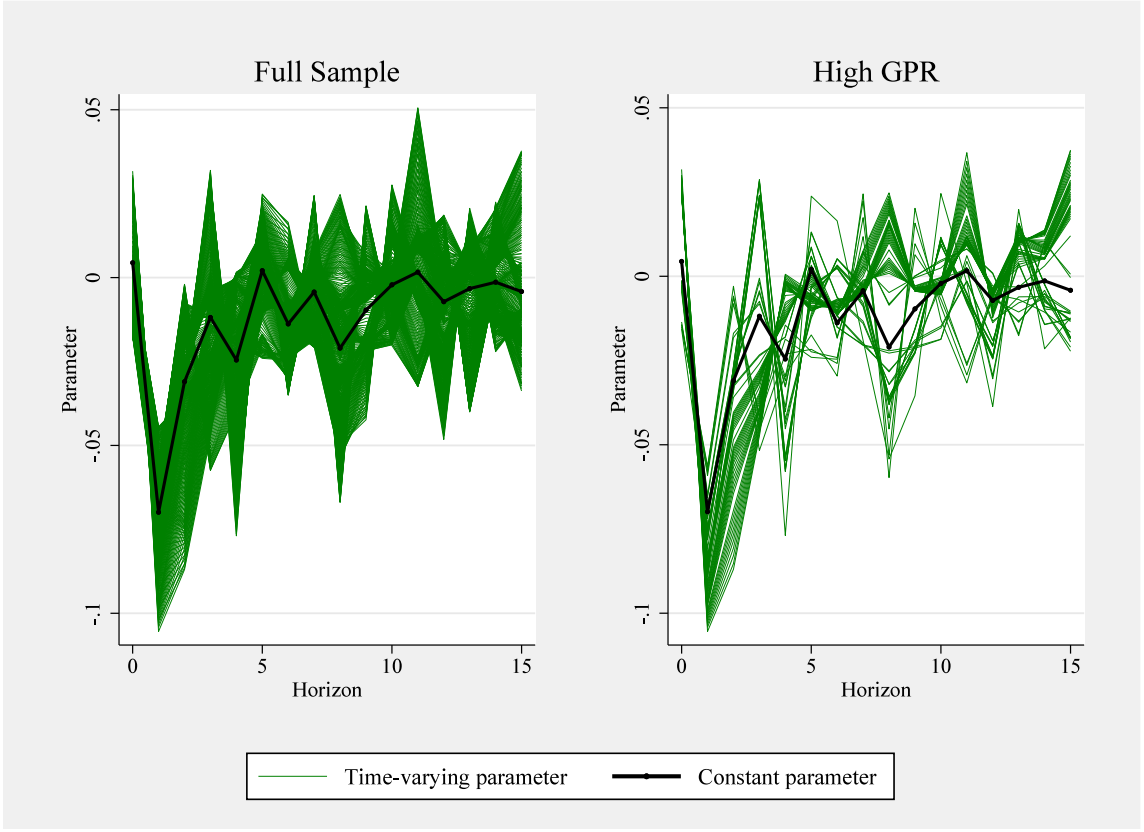


Figure A.18: Time-varying responses of the median-minus-95th-percentile gap.

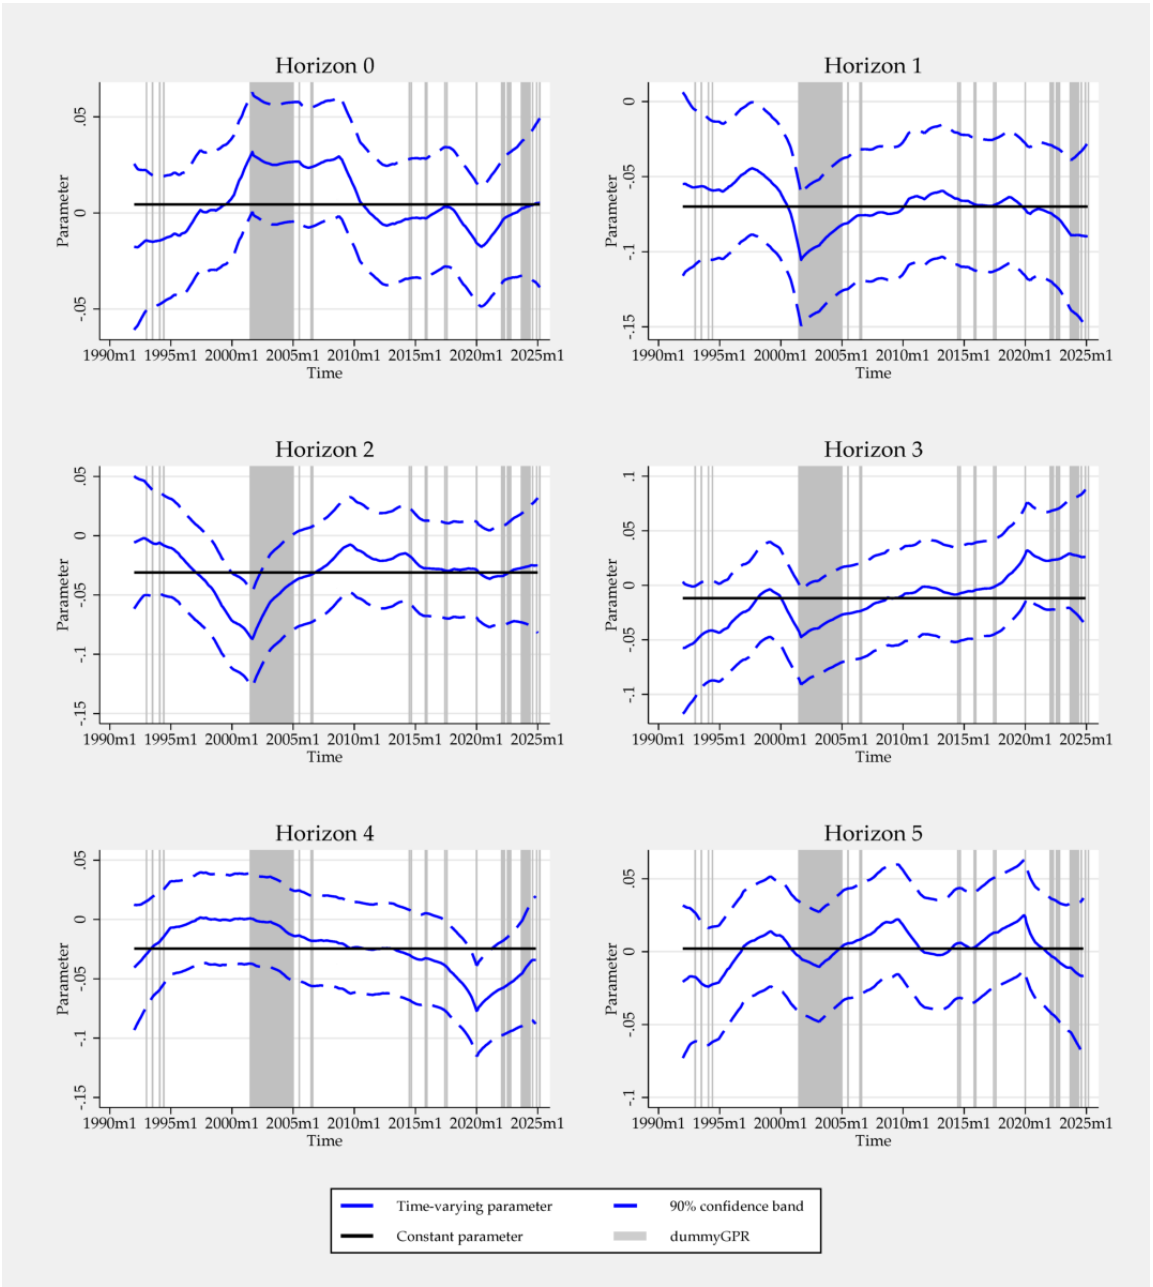


Figure A.19: Median-minus-95th-percentile response across time at horizons zero through five. Dashed lines are 90% confidence bands.

A.6 Alternative estimator

As an alternative to conditional quantiles, we estimate distribution regressions (Foresi and Peracchi, 1995). For each threshold k, define $d_{t+1,k} = 1\{y_{t+1} \leq k\}$ and estimate a logit of $d_{t+1,k}$ on GPR and current investment growth. A positive GPR coefficient means that higher GPR is associated with a greater probability of next-quarter growth falling below k. Figure A.20 shows positive and statistically meaningful coefficients at severe negative thresholds, whereas estimates around favorable thresholds are small and imprecise. This corroborates the direction of the baseline result without imposing a linear conditional-quantile function.

A.7 Supplementary outcomes

The remaining exercises are mechanism-consistent patterns, not mechanism identification. The lower-tail GPR coefficient is larger for durable than for nondurable consumption, while the association for services consumption is weak (Figures A.21–A.23). This ranking is consistent with postponement when purchases are durable. The absence of a precise association with aggregate corporate bond yields (Figure A.24) does not rule out heterogeneous financing channels. Finally, GPR does not display the same quantile pattern for next-quarter GDP growth in this U.S. time-series specification (Figure A.25). This last result concerns the present model and should not be read as contradicting the broader cross-country evidence of Reynolds and Lloyd (2026).

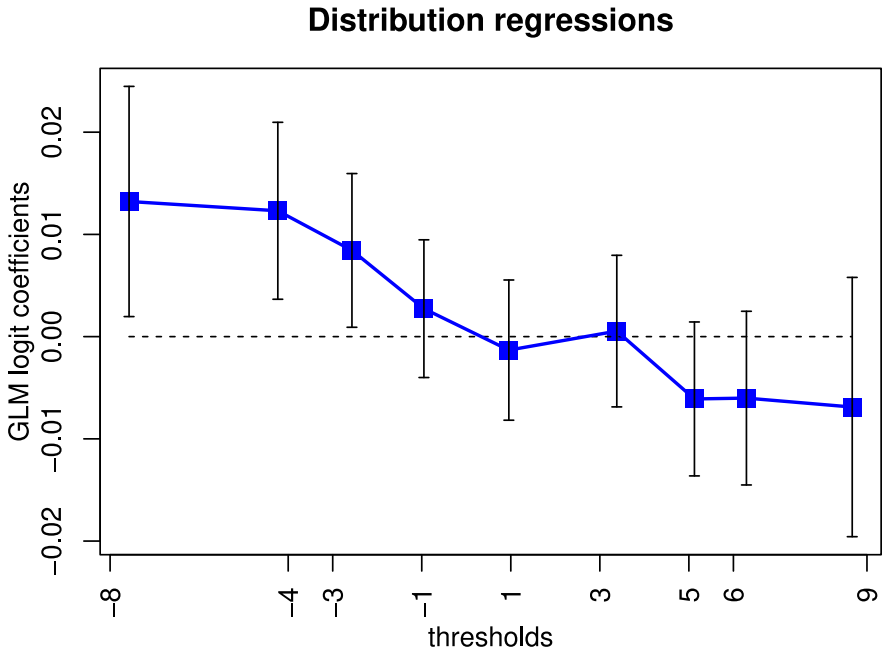


Figure A.20: Distribution-regression coefficients. Whiskers are 90% HAC confidence intervals.

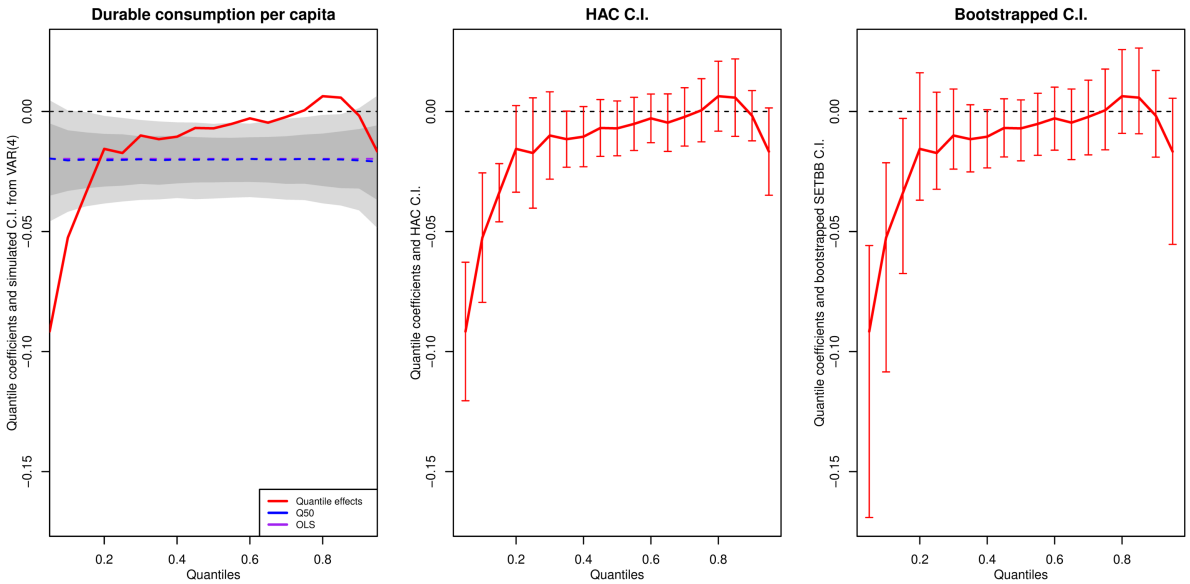


Figure A.21: GPR and the conditional quantiles of next-quarter durable-consumption growth.

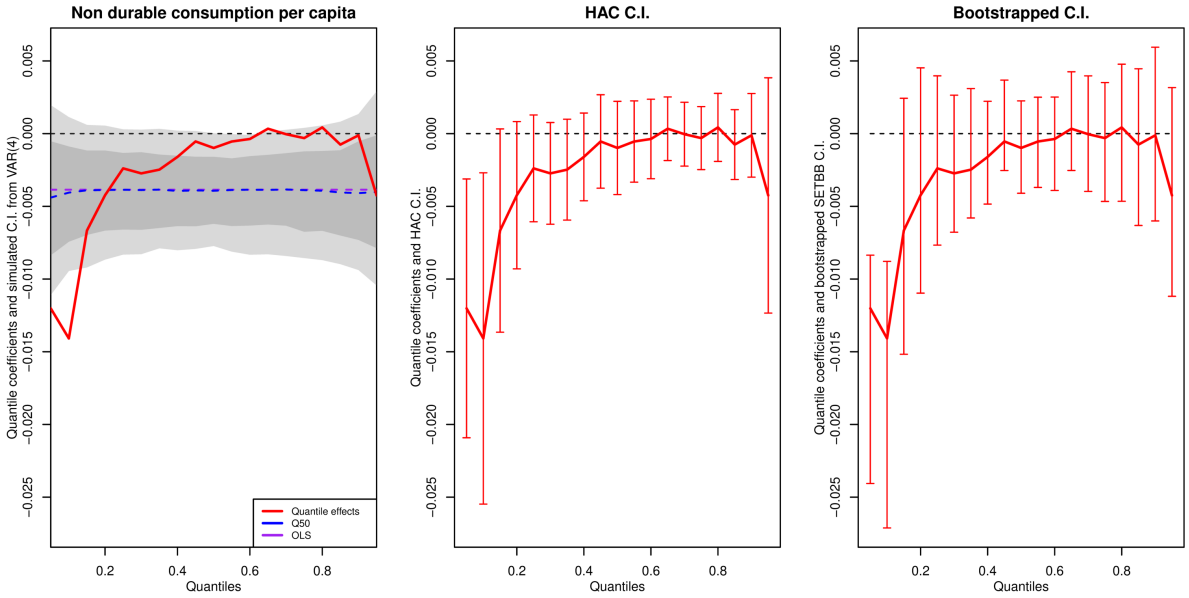


Figure A.22: GPR and the conditional quantiles of next-quarter nondurable-consumption growth.

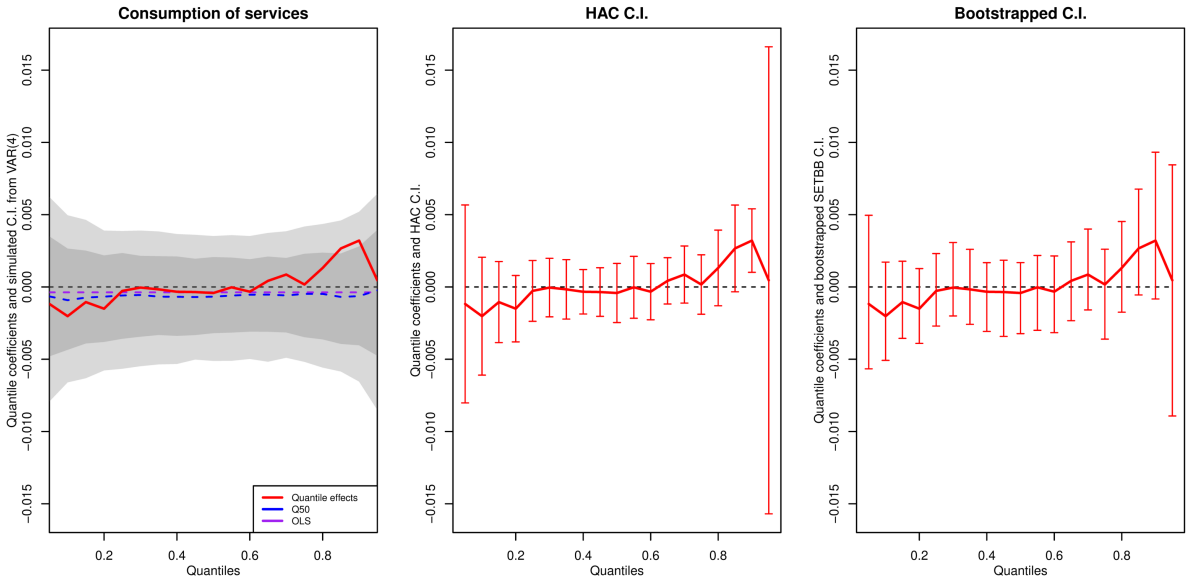


Figure A.23: GPR and the conditional quantiles of next-quarter services-consumption growth.

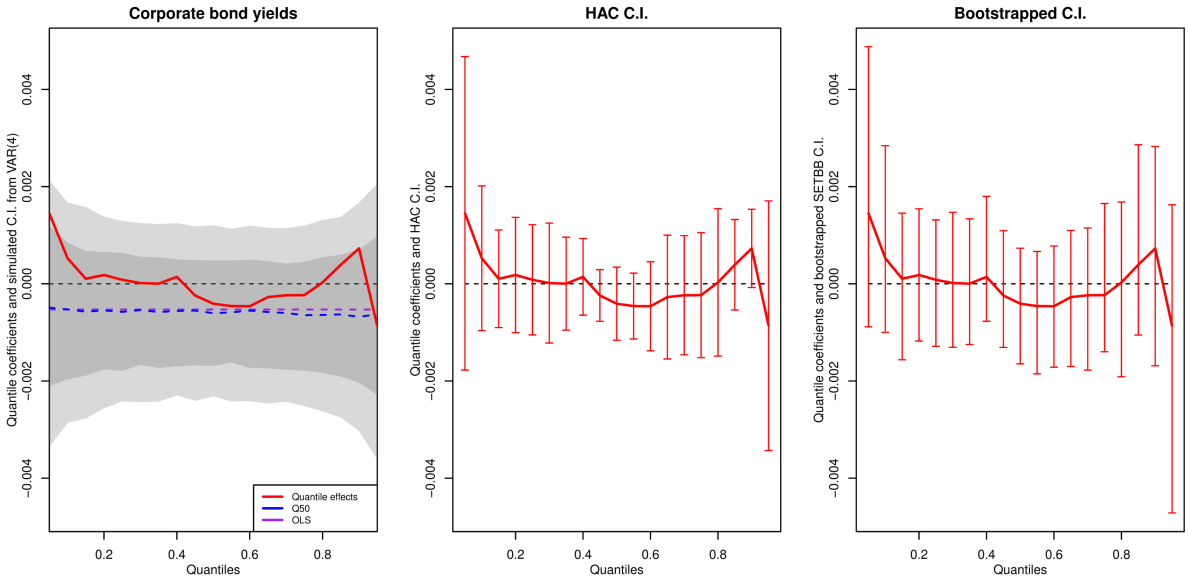


Figure A.24: GPR and the conditional quantiles of next-quarter corporate bond yields.

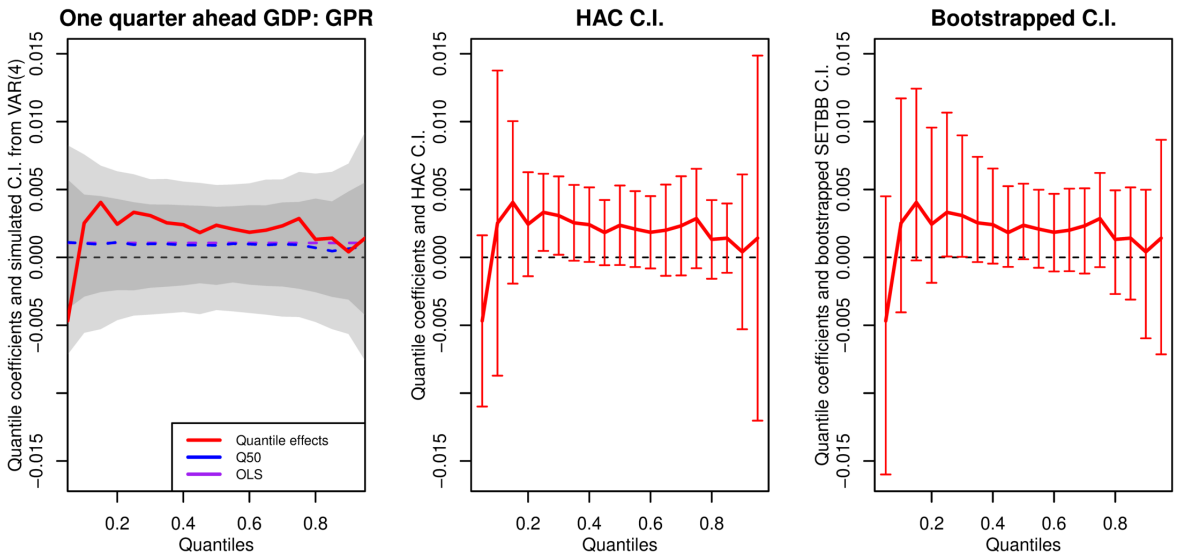


Figure A.25: GPR and the conditional quantiles of next-quarter real GDP growth.

A.8 Formal comparison of threats and acts

To compare the two components, we estimate separate threats and acts versions of equation (1) using the same outcome, sample, and controls. At each quantile $\tau = 0.05, 0.10, \ldots, 0.95$, the null is $H_0 : \beta_\tau^T - \beta_\tau^A = 0$. Define the estimated difference and its bootstrap pivot as

$$\widehat{d}_\tau = \widehat{\beta}_\tau^T - \widehat{\beta}_\tau^A, \qquad P_{\tau,d}^{(b)} = P_{\tau,T}^{(b)} - P_{\tau,A}^{(b)}, \quad b = 1, \ldots, 1000. \qquad (A3)$$

The smooth extended tapered block bootstrap of Gregory et al. (2018), implemented in QregBB, supplies the coefficient pivots. Matched SETBB replications pair the threats and acts specifications on the common sample. The paired construction retains the covariance of the coefficient estimates, rather than assuming their estimation errors are independent. Using empirical quantiles $c_{a,\tau}$ of the difference pivots, the 90% interval is

$$\left[\widehat{d}_\tau - \frac{c_{0.95,\tau}}{\sqrt{n}},\ \widehat{d}_\tau - \frac{c_{0.05,\tau}}{\sqrt{n}}\right], \qquad (A4)$$

where n is the common estimation-sample size. The returned SETBB pivots already incorporate their centering and taper adjustment. The comparison is in the index units used in the separate regressions and concerns equally sized changes in those regressors.

Figure A.26 shows negative estimated differences in the lower tail, including the fifth percentile. All 90% pointwise intervals nevertheless contain zero. Equality cannot be rejected at the 10% two-sided level; this is neither evidence of equality nor a simultaneous test across quantiles. The point estimates are consistent with greater lower-tail sensitivity to threats, but the comparison does not establish a statistically significant ranking or uniquely identify an option-value-of-waiting mechanism.

We also test the joint lower-tail null $H_0 : d_{0.05} = d_{0.10} = 0$, where $d_\tau = \beta_\tau^T - \beta_\tau^A$. Let $\widehat{\mathbf{d}} = (\widehat{d}_{0.05}, \widehat{d}_{0.10})'$. The Wald statistic is $W = \widehat{\mathbf{d}}' \widehat{V}^{-1} \widehat{\mathbf{d}}$, where $\widehat{V}$ is the covariance matrix of the two estimated differences obtained from the joint distribution of the corresponding SETBB difference pivots. We obtain $W = 0.390$, with a bootstrap-calibrated p-value of 0.843 based on the centered block-bootstrap distribution of the statistic and an asymptotic $\chi^2(2)$ p-value of 0.823. The joint null is not rejected, consistent with the pointwise intervals.

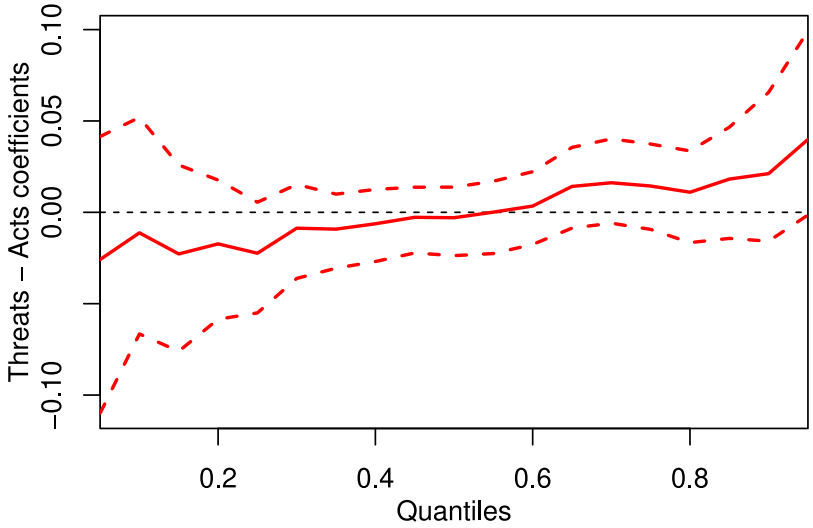


Figure A.26: Difference between the quantile coefficients on geopolitical threats and acts. The solid line reports $\widehat{\beta}_\tau^T - \widehat{\beta}_\tau^A$; dashed lines are 90% pointwise intervals from 1,000 paired SETBB replications. Negative values indicate a more negative threats coefficient. The horizontal dashed line denotes zero.

References

Barrett, P. (2025). Estimated monthly national accounts for the United States. *IMF Working Papers*, 2025(134).

Caldara, D. and Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4):1194–1225.

Foresi, S. and Peracchi, F. (1995). The conditional distribution of excess returns: An empirical analysis. *Journal of the American Statistical Association*, 90(430):451–466.

Gregory, K. B., Lahiri, S. N., and Nordman, D. J. (2018). A smooth block bootstrap for quantile regression with time series. *Annals of Statistics*, 46(3):1138–1166.

Inoue, A., Rossi, B., and Wang, Y. (2024). Local projections in unstable environments. *Journal of Econometrics*, 244(2):105726.

Loria, F., Matthes, C., and Zhang, D. (2025). Assessing macroeconomic tail risk. *Economic Journal*, 135(665):264–284.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local projections and VARs estimate the same impulse responses. *Econometrica*, 89(2):955–980.

Püttmann, L. (2018). Patterns of panic: Financial crisis language in historical newspapers. Available at SSRN 3156287.

Reynolds, J. and Lloyd, S. (2026). Tension across the tail: The cross-country effects of geopolitical risk on macroeconomic tail risk. *Economics Letters*, 264:112927.