Conditional Excess Volatility

CHRISTIAN SKOV JENSEN
Copenhagen & Bocconi

EBEN LAZARUS
UC Berkeley Haas & NBER

NOVEMBER 2026

Overview

Market volatility rises significantly in bad times. How much of it is compensated?

- **Standard theory:** Risk premium rises with market variance. E.g., with a constant price of risk:

$$\underset{\text{[risk premium]}}{\mu_t} = \underset{\text{[price of risk]}}{\widetilde{\gamma}} \times \underset{\text{[market variance]}}{\sigma_t^2}$$

- **2008 Financial Crisis:** Market volatility rose from 15% to 83% at the height of the crisis...

$$\widetilde{\gamma} = 1 \implies \mu_t = \text{ 69\% per year}$$
$$\widetilde{\gamma} = 3 \implies \mu_t = 207\% \text{ per year}$$

...far above plausible estimates of the risk premium

$\implies$ **Much of the increase in σ_t^2 must be uncompensated, so "too much" price variation**

[and this applies for general notions of $\widetilde{\gamma}$]

Risk Prices in Bad Times

Price per unit of variance risk pins down exposure of marginal utility to the market:

$$\gamma_t \equiv \frac{\mu_t}{\sigma_t^2} = -R_{f,t+1}\,\beta_{M\to R,t}$$

[observed] [beta of SDF on market]

- γ_t declines in bad times: Sharpe ratio $\frac{\mu_t}{\sigma_t}$ tends to rise, but volatility σ_t rises by more
- **Puzzle:** Why is the SDF *less* exposed to the market when volatility & risk aversion seem highest?

Risk Prices in Bad Times

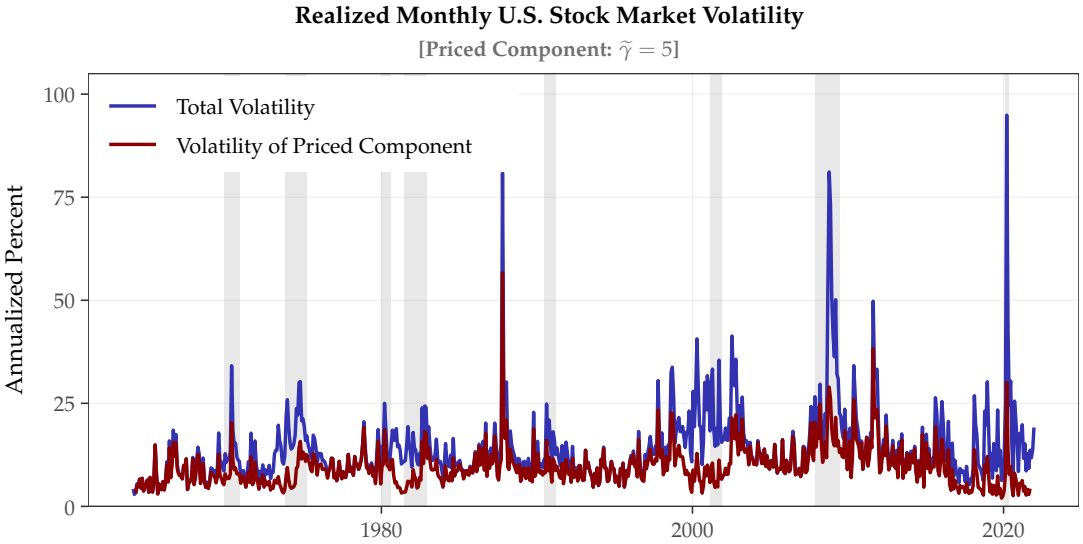
General decomposition:

$$\gamma_t \equiv \frac{\mu_t}{\sigma_t^2} = \widetilde{\gamma}_t \times (1 - \psi_t)$$

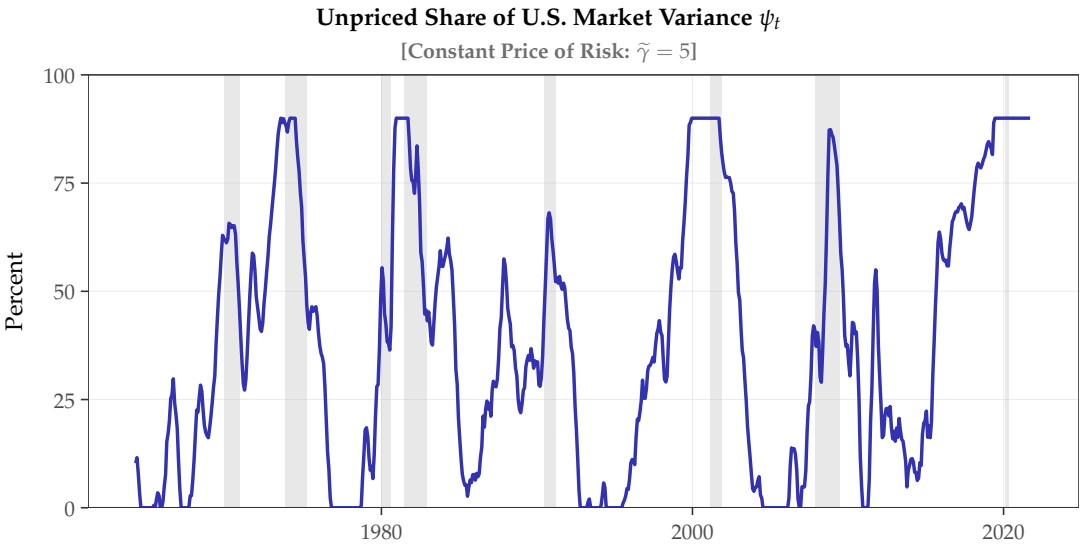
[observed] [true price of risk] [priced share of variance]

- γ_t declines in bad times: Sharpe ratio $\frac{\mu_t}{\sigma_t}$ tends to rise, but volatility σ_t rises by more
- **Puzzle:** Why is the SDF *less* exposed to the market when volatility & risk aversion seem highest?
- Two possibilities: True price of risk $\widetilde{\gamma}_t$ falls in bad times [unlikely], or unpriced share of variance ψ_t rises

Priced and Unpriced Market Volatility



Priced and Unpriced Market Volatility



Conditional Excess Volatility

"Classic" excess volatility [Shiller 1981]**:** Prices move far more than fundamental value...

- ...but need a model of discount rates to measure fundamental value

Conditional version, in terms of risk premium & variance:

- Use good times as a benchmark
- To rationalize higher variance in bad times, need risk prem. $\nearrow$ at least in proportion to variance
- No measure of fundamental value needed
  [instead: weak assumption on cyclicality of true price of risk]

$\Longrightarrow$ **Lower bound on excess *(unpriced)* volatility in bad times**

Main Results

1. **Theory:** Simple characterization of $\gamma_t = \frac{\mu_t}{\sigma_t^2} \Longrightarrow$ lower bound on excess vol. in bad times

2. **Excess volatility:** Unpriced vol share $\psi_t \geqslant 50\%$ in bad times (under moderate incr. in true risk price $\widetilde{\gamma}_t$)
   - **Measure** γ_t (realized, expected, & option-implied) & **put structure on** $\widetilde{\gamma}_t$ to **infer** ψ_t
   - γ_t tends to fall in recessions, in the U.S. and across 20 stock markets
   - Options: Risk price falls with horizon, as implied by procyclical γ_t

3. **Implications:**
   - Portfolio choice: Optimal market weight doesn't increase in bad times despite higher Sharpe ratios
   - Real investment: High excess volatility predicts lower investment, despite being unpriced

4. **Model:** Unpriced variance explains downward-sloping term structure of dividend claim Sharpe ratios

Related Literature

1. Excess volatility [Shiller 1981; LeRoy & Porter 1981; Campbell & Shiller 1987]
   - Conditional version without knowledge of fundamental value: Lower bound on excess volatility in bad times
2. Option-implied risk aversion [Aït-Sahalia & Lo 2000; Jackwerth 2000; Bliss & Panigirtzoglou 2004]
   - γ_t as a measure of risk aversion understates true price of risk, especially in bad times
3. Risk & return over the cycle [Lustig & Verdelhan 2012; Moreira & Muir 2017; Schreindorfer & Sichert 2023]
   - Unpriced variance rises in bad times, which helps explain why volatility-managed portfolios earn high alphas
4. Uncertainty & investment [Bernanke 1983; Bloom 2009]
   - Unpriced component of market volatility predicts lower investment
5. Term structure of equity risk premia [Lettau & Wachter 2007; van Binsbergen & Koijen 2017; Gormsen 2021]
   - Sharpe ratios are higher on claims to near-term dividends because unpriced share rises with maturity

Roadmap

1. Introduction

2. Theory: Priced & Unpriced Variance
   - Decomposition
   - Standard Models

3. Excess Volatility in Bad Times

4. Implications & A Simple Model

5. Final Notes

Decomposing Market Variance

Given $\mathbb{E}_t[M_{t+1}R_{t+1}] = 1$ for SDF M_{t+1}, market risk premium is

$$\mu_t \equiv \mathbb{E}_t[R_{m,t+1} - R_{f,t+1}] = -R_{f,t+1}\,\mathrm{Cov}_t(M_{t+1}, R_{m,t+1})$$

Decompose mkt return: $R_{m,t+1} - \mathbb{E}_t[R_{m,t+1}] = s_{t+1} + u_{t+1}, \quad s_{t+1} \propto (M_{t+1} - \mathbb{E}_t[M_{t+1}]), \quad u_{t+1} \perp M_{t+1}$

$$\sigma_t^2 \quad = \quad \underbrace{\mathrm{Var}_t(s_{t+1})}_{\sigma^2_{S,t}} \quad + \quad \underbrace{\mathrm{Var}_t(u_{t+1})}_{\sigma^2_{\perp,t}}$$

[spanned: priced] [orthogonal: unpriced]

- Only spanned variance commands a risk premium: $\mu_t = -R_{f,t+1}\,\mathrm{Cov}_t(M_{t+1}, s_{t+1})$
- **Conditional excess volatility** $\sigma^2_{\perp,t}$: Return variance not compensated by a risk premium
- Market return spanned by the SDF $\Longrightarrow \sigma^2_{\perp,t} = 0$, while non-fundamental shocks $\Longrightarrow \sigma^2_{\perp,t} > 0$

Decomposing Market Variance

Given $\mathbb{E}_t[M_{t+1}R_{t+1}] = 1$ for SDF M_{t+1}, market risk premium is

$$\mu_t \equiv \mathbb{E}_t[R_{m,t+1} - R_{f,t+1}] = -R_{f,t+1}\,\mathrm{Cov}_t(M_{t+1}, R_{m,t+1})$$

Decompose mkt return: $R_{m,t+1} - \mathbb{E}_t[R_{m,t+1}] = s_{t+1} + u_{t+1}$, $s_{t+1} \propto (M_{t+1} - \mathbb{E}_t[M_{t+1}])$, $u_{t+1} \perp M_{t+1}$

$$\sigma_t^2 = \underbrace{\mathrm{Var}_t(s_{t+1})}_{\sigma^2_{S,t}} + \underbrace{\mathrm{Var}_t(u_{t+1})}_{\sigma^2_{\perp,t}}$$

[spanned: priced] [orthogonal: unpriced]

- Non-fundamental shocks $\Longrightarrow \sigma^2_{\perp,t} > 0$
- Also: priced risks not spanned by the market (e.g., labor income) $\Longrightarrow \sigma^2_{\perp,t} > 0$
- . . .**but** priced risks not spanned by the market $\Longleftrightarrow$ unpriced component of mkt returns

Details

Measuring Excess Volatility from γ_t

Excess volatility share:	ψ_t	$= \sigma^2_{\perp,t}/\sigma^2_t$	
Price per unit of spanned variance:	$\widetilde{\gamma}_t$	$= \mu_t/\sigma^2_{S,t}$	[true price of risk]
Price per unit of variance risk:	γ_t	$= \mu_t/\sigma^2_t$	[observable mean-var. ratio]

Result 1 (*Decomposition & Interpretation*)

The price per unit of variance risk satisfies

$$\gamma_t = \widetilde{\gamma}_t\,(1-\psi_t) = -R_{f,t+1}\underbrace{\beta_{M\to R,t}}_{\mathrm{Cov}_t(M_{t+1},R_{m,t+1})/\sigma^2_t}$$

- γ_t understates **true price of risk** $\widetilde{\gamma}_t$: Some of the variance in the denominator is uncompensated
- With constant $\widetilde{\gamma}$, time variation in γ_t reflects time variation in ψ_t: $\psi_t = 1-\gamma_t/\widetilde{\gamma}$
- Unpriced noise in $R_{m,t+1}$ attenuates $\beta_{M\to R,t} \implies$ SDF less exposed to market even if $\widetilde{\gamma}_t$ ↗

Excess Volatility in Canonical Models

Excess Volatility Across Models: Simulations for Monthly Log Returns

[Good & Bad Times: Terciles of Main State Variable]

	γ_t		Priced $\widetilde{\gamma}_t$		Ex. vol. share ψ_t (%)	
	Good	Bad	Good	Bad	Good	Bad
CRRA, i.i.d. consumption	2.0	2.0	2.0	2.0	0.0	0.0
Campbell–Cochrane	2.0	4.1	2.0	4.1	0.0	0.1
Bansal–Yaron	3.8	3.9	3.9	4.0	2.0	1.7
Drechsler–Yaron	2.9	3.1	3.1	3.5	5.3	11
Wachter (disasters)	1.2	1.3	1.2	1.3	0.1	1.1
Lettau–Wachter	0.4	4.7	3.9	33.8	88	86

- With priced dividend shocks, market is close to spanning the SDF: ψ_t is small Details
- Lettau–Wachter: Adding unpriced shocks to price of risk generates high ψ_t & the value premium
- But all miss the cyclicality: Weakly countercyclical γ_t & $\widetilde{\gamma}_t$, and small or flat excess vol.
- **Empirically:** Ratio of γ_t in very bad vs. normal times ≈ 0.70

Why Excess Volatility Must Rise in Crises

Good times G **& crisis** B**:** True price of risk $\widetilde{\gamma}_B \geqslant \widetilde{\gamma}_G$ in standard models

Result 2 (*Lower Bound*)

If $\widetilde{\gamma}_B \geqslant \widetilde{\gamma}_G$, this implies a lower bound on the share of excess volatility in bad times:

$$\psi_B \geqslant 1 - \frac{\gamma_B}{\gamma_G}(1 - \psi_G)$$

If the observed $\gamma_t = \mu_t/\sigma_t^2$ declines in bad times ($\gamma_B \leqslant \gamma_G$), this lower bound is positive.

- Most conservative case, $\psi_G = 0$: $\psi_B \geqslant 1 - \gamma_B/\gamma_G$
- γ_t halves in crises $\Longrightarrow$ at least half of crisis-time variance is excess volatility
- If γ_t is flat, ψ_t must rise with $\widetilde{\gamma}_t$: $\widetilde{\gamma}_B = 2\,\widetilde{\gamma}_G$ & $\psi_G = 0.2$ imply $\psi_B = 0.6$

Why Excess Volatility Must Rise in Crises

Good times G **& crisis** B**:** True price of risk $\widetilde{\gamma}_B \geqslant \widetilde{\gamma}_G$ in standard models

Result 2 (*Lower Bound*)

If $\widetilde{\gamma}_B \geqslant \widetilde{\gamma}_G$, this implies a lower bound on the share of excess volatility in bad times:

$$\psi_B \geqslant 1 - \frac{\gamma_B}{\gamma_G}(1 - \psi_G)$$

If the observed $\gamma_t = \mu_t / \sigma_t^2$ declines in bad times ($\gamma_B \leqslant \gamma_G$), this lower bound is positive.

- Most conservative case, $\psi_G = 0$: $\psi_B \geqslant 1 - \gamma_B / \gamma_G$
- Recall $\gamma_t \approx -\beta_{M \to R,t}$
- If risk aversion ↗ in crises, can rationalize lower exposure of MU to market only with excess vol. ↗

Roadmap

1. Introduction

2. Theory: Priced & Unpriced Variance

3. Excess Volatility in Bad Times
    - Measuring γ_t
    - Cyclicality
    - Option Evidence

4. Implications & A Simple Model

5. Final Notes

Three Measures of γ_t

Goal: Mean-variance ratio $\gamma_t = \mu_t / \sigma_t^2$ each month, to compare good & bad times

$\Longrightarrow$ **lower bound on excess volatility in bad times**

1. **Realized:** Mean over variance of daily excess returns within each month
   [U.S. since 1926]

2. **Expected:** Risk premium forecast from the cross-section of book-to-market ratios, divided by AR(1) variance forecast [U.S. since 1964, based on Kelly & Pruitt 2013]
   - $\mu_t = \alpha + \beta \sum_i w_i \, bm_{i,t}$ over 100 size & book-to-market portfolios, with $w_i \propto \text{Cov}(bm_{i,t}, R_{m,t+1})$
   - Ann. $R^2 = 18\%$ in sample, 13% OOS (higher weights: large firms, growth firms)
   - Results hold or strengthen with ML return & HAR or EGARCH variance forecasts Alt. forecasts

Three Measures of γ_t

Goal: Mean-variance ratio $\gamma_t = \mu_t / \sigma_t^2$ each month, to compare good & bad times

$\Longrightarrow$ **lower bound on excess volatility in bad times**

1. **Realized:** Mean over variance of daily excess returns within each month
   [U.S. since 1926]

2. **Expected:** Risk premium forecast from the cross-section of book-to-market ratios, divided by AR(1) variance forecast [U.S. since 1964, based on Kelly & Pruitt 2013]

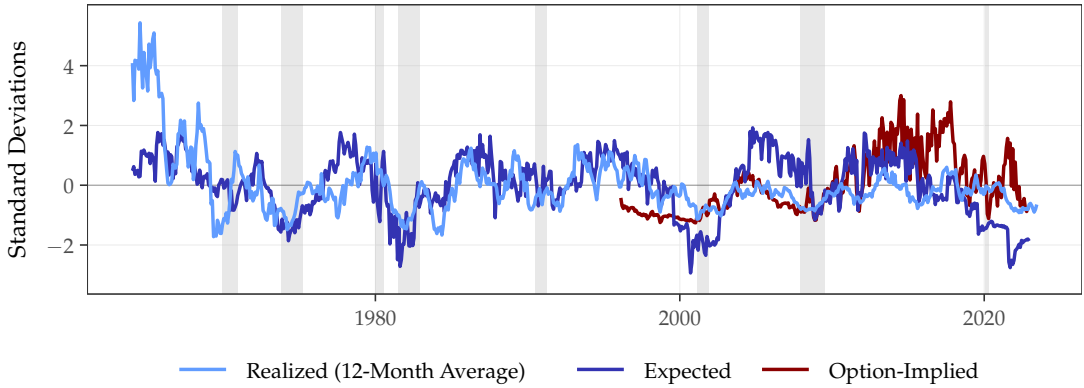
3. **Option-implied:** Compare S&P options to realized returns to estimate γ_t in $M_{t+1} \propto R_{m,t+1}^{-\gamma_t}$
   [U.S. since 1996, based on Bliss & Panigirtzoglou 2004]

Details

International: Realized & expected measures across 20 stock markets

Estimated Mean-Variance Ratio

Standardized U.S. price per unit of variance risk γ_t:



- Despite different data & methods, the three measures comove [correlations 0.29–0.42]
- γ_t typically lower in recessions

Sharpe ratios

Mean-Variance Ratio in Recessions

Recession Minus Non-Recession Months

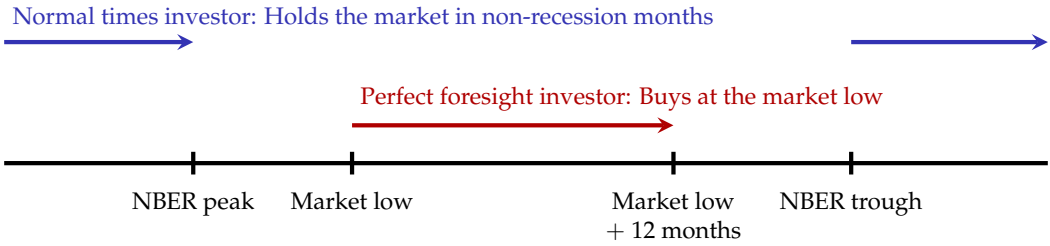
	Price per unit of var. risk, $\gamma_B - \gamma_G$	
	Post-1964	Post-1996
Realized	-22.35^{***}	-16.37^{***}
	(5.32)	(2.90)
Expected	-1.85^{***}	-1.29
	(0.49)	(0.86)
Option-implied		-1.17^{***}
		(0.33)

Standard errors are HAR [Lazarus et al. (2018)].

- Start by conditioning on macro state *(binary, not directly related to σ_t)*
- Price of variance risk γ_t much lower in recessions
- **Problem:** Recessions are dated ex post, and stock prices often fall early in them *(a form of selection bias)*

Correcting for Selection Bias

Instead, exclude recession months before the market low. Compare two investors:



- Perfect foresight avoids attributing the downturn to a lower ex ante expected return
  (best chance for countercyclical γ_t)
- If risk prices do rise, perfect foresight investor should earn a higher price per unit of variance risk
- Compare price per unit of variance risk γ_t & Sharpe ratios for both investors

Mean-Variance vs. Sharpe Ratio After Recession Lows

Perfect-Foresight Recession Investor Minus Normal-Times Investor

	Price per unit of var. risk, $\gamma_B - \gamma_G$		Sharpe ratio, $SR_B - SR_G$	
	Post-1964	Post-1996	Post-1964	Post-1996
Realized	−1.79	−4.78	1.13**	0.43
	(6.70)	(4.97)	(0.47)	(0.85)
Expected	−1.77***	−2.23***	−0.17*	−0.26
	(0.41)	(0.66)	(0.08)	(0.16)
Option-implied		−0.79**		0.10*
		(0.34)		(0.05)

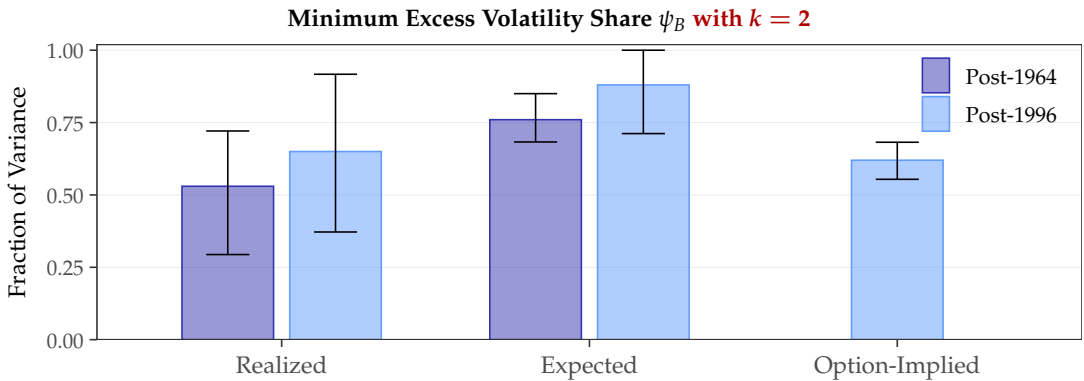
HAR standard errors in parentheses.

- Price of variance risk γ_t lower for all three measures
- Sharpe ratios higher after market lows for realized & option-implied measures
- Excess volatility share ψ_t rises in bad times

Recent crises

How Much Excess Volatility in Recessions?

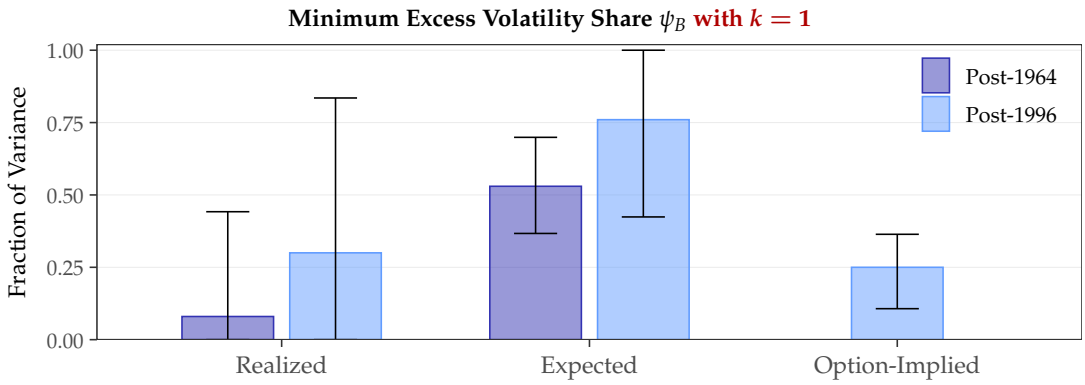
If $\widetilde{\gamma}_B = k\,\widetilde{\gamma}_G$ in recessions: $\psi_B \geqslant 1 - \gamma_B/(k\,\gamma_G)$



- If $k = 2$, more than half of market variance is unpriced
  [Campbell–Cochrane: $k \approx 2$]

How Much Excess Volatility in Recessions?

If $\widetilde{\gamma}_B = k\,\widetilde{\gamma}_G$ in recessions: $\psi_B \geqslant 1 - \gamma_B / (k\,\gamma_G)$



- Even with $k = 1$, still have meaningful conditional excess volatility in recessions
  [most conservative case: constant true price of risk]

International Evidence

Is this specific to the U.S.? Perfect-foresight investor in 20 stock markets, with OECD recession dates:

Recession Minus Non-Recession

	γ^{realized}	γ^{expected}	Sharpe ratio
Recession − non-recession	0.99	−0.41	0.60**
	(1.54)	(0.33)	(0.24)
γ_B/γ_G	1.08	0.88	
Minimum ψ_B $(k = 2)$	0.46	0.56	

Standard errors clustered by index & date.

- Price per unit of variance risk γ_t doesn't rise in recessions
- Sharpe ratio higher in recessions
- With $k = 2$: Minimum excess volatility share ψ_B of 0.46 for realized & 0.56 for expected γ_t

Mean-Variance Ratio and Financial & Macro Conditions

Regressions of U.S. γ_t & Sharpe Ratio on Financial & Macro Indicators

	Price per unit of variance risk			Sharpe ratio		
	γ^{realized}	γ^{expected}	γ^{option}	SR^{realized}	SR^{expected}	SR^{option}
NFCI	−5.47**	−0.65***	−0.93**	−0.16***	−0.02*	0.01
	(2.30)	(0.20)	(0.37)	(0.05)	(0.01)	(0.01)
Recession prob.	−22.71***	−2.18***	−1.58***	−0.46***	−0.07*	0.03
	(5.80)	(0.69)	(0.47)	(0.14)	(0.04)	(0.02)

HAR standard errors in parentheses.

- Price of variance risk γ_t falls when financial conditions tighten & recession probability rises
- Sharpe ratios less sensitive to these variables
- International panel: γ_t rises with future consumption & industrial production growth
- Consumption-based γ_t falls in recessions: Market vol. ↗ more than cons. vol.

Consumption

Option Evidence: Risk Prices by Horizon

Independent out-of-sample test of procyclical γ_t, from option portfolios

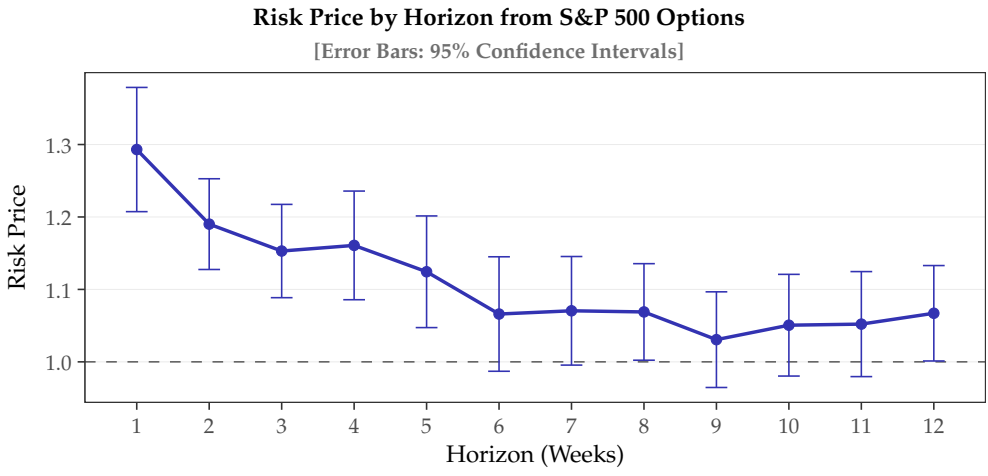
Binary bet: Short an Arrow–Debreu claim on return state ω_1, long one on $\omega_2 > \omega_1$ (return $R_{\omega,t+1}$)

- Bet fixes the quantity of risk: Expected return depends only on the price per unit of variance risk γ_t

$$\log \mathbb{E}_t[R_{\omega,t+1}] = \gamma_t \left(\omega_2 - \omega_1\right)$$

- Lower $\gamma_t \implies$ higher price on a multi-period bet
- Procyclical price of variance risk $\gamma_t \iff$ multi-period bets hedge against bad states, lower expected returns at longer horizons
- Test: Risk price $\mathbb{E}_t[R_{\omega,t+1}]$ by horizon from S&P 500 options, 1996–2018

Option Evidence: Risk Prices by Horizon



- Decreases with horizon, consistent with procyclical price of variance risk γ_t, ex. vol. in bad times

Roadmap

1. Introduction

2. Theory: Priced & Unpriced Variance

3. Excess Volatility in Bad Times

4. Implications & A Simple Model
    - Portfolio Choice & Investment
    - Equity Term Structure Model

5. Final Notes

Portfolio Choice and Volatility Timing

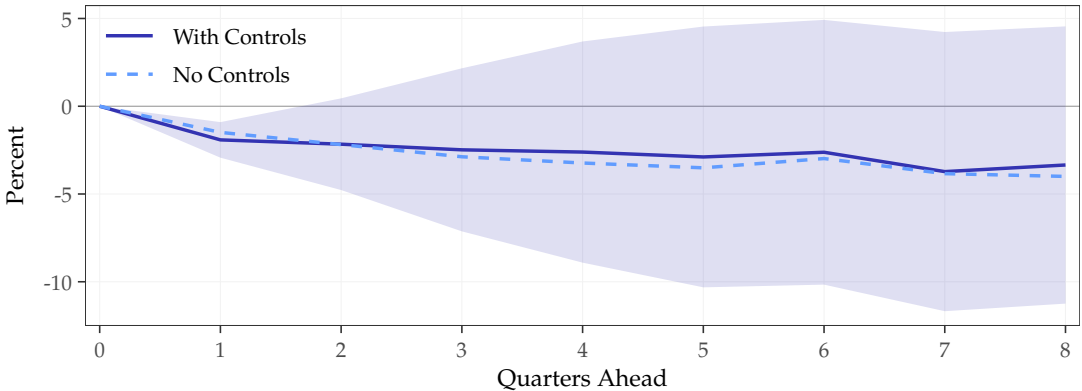
Myopic investor with relative risk aversion $\widetilde{\gamma}$ equal to true price of risk:

$$\underset{\text{[portfolio weight]}}{w_{m,t}} = \frac{\mu_t}{\widetilde{\gamma}\,\sigma_t^2} = \frac{\gamma_t}{\widetilde{\gamma}} = 1 - \psi_t$$

- Mean-variance ratio γ_t as measure of risk aversion: Understates $\widetilde{\gamma}_t$, especially in bad times
- Optimal portfolio weight $\searrow$ in bad times, even as Sharpe ratio $\nearrow$: Excess volatility share ψ_t $\nearrow$
  [standard theory: $\psi_t = 0 \Longrightarrow w_{m,t} = 1$]
- Scaling down in high-variance months reduces exposure to unpriced risk, which helps explain why volatility-managed portfolios earn alpha [Moreira & Muir 2017]

Excess Volatility and Real Investment

Real effects of unpriced risk? Investment growth after high ψ_t (expected γ_t below 25th pctile)**:**

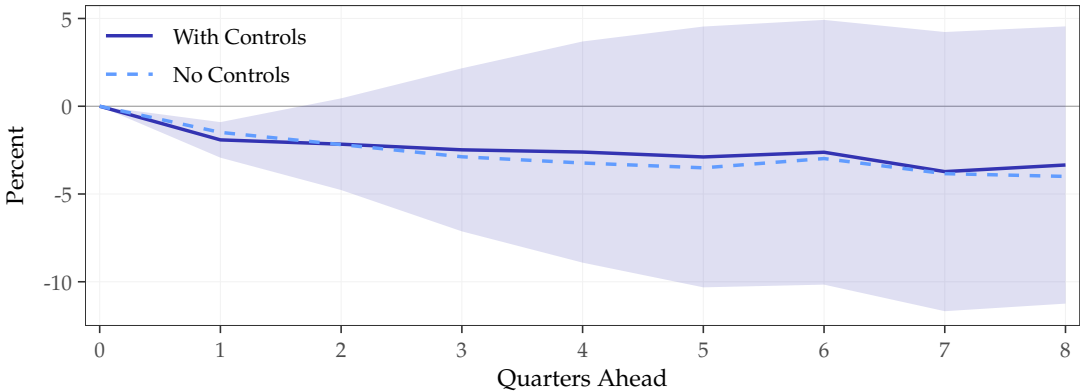


- Lower investment growth over the next one to two quarters, with and without controls
- Controls: VIX, rec. prob., NFCI, div. yield, credit spread

Details

Excess Volatility and Real Investment

Real effects of unpriced risk? Investment growth after high ψ_t (expected γ_t below 25th pctile)**:**



- Lower investment growth over the next one to two quarters, with and without controls
- Apparent real effects from ex. vol. despite being unpriced

Details

A Simple Model of Conditional Excess Volatility

- Builds on Lettau & Wachter (2007): Price of risk x_t rises in bad times, and dividend shocks are priced
- **Additional unpriced dividend shock** with volatility increasing in x_t
- How model matches the evidence:
  - Procyclical γ_t: Risk premium rises with x_t, but unpriced variance rises by more
  - Sharpe ratios are countercyclical
  - Sharpe ratios are higher on claims to near-term dividends [van Binsbergen & Koijen 2017], because longer maturities more exposed to unpriced x_t shocks *(creates vol. w/ no compensation)*

$\Longrightarrow$ unpriced excess vol. in bad times helps make sense of multiple asset pricing facts

Formal Setup

- Price of risk x_t follows an AR(1), and dividend shocks are priced:

$$m_{t+1} = -r^f - \tfrac{1}{2}x_t^2 - x_t\,\varepsilon_{d,t+1}$$

- Unpriced dividend shock $\varepsilon_{z,t+1}$ with volatility $x_t\sigma_z$:

$$\Delta d_{t+1} = g - \tfrac{1}{2}x_t^2\sigma_z^2 + \sigma_d\,\varepsilon_{d,t+1} + x_t\sigma_z\,\varepsilon_{z,t+1}$$

- Price per unit of variance risk for the n-maturity dividend claim:

$$\gamma_{n,t} = \frac{\sigma_d\,x_t}{\sigma_d^2 + \sigma_z^2 x_t^2 + B_{x,n-1}^2\sigma_x^2}$$

[σ_x: volatility of shocks to x_t; $B_{x,n-1}$: loading of log p/d ratio on x_t]

- Large unpriced dividend volatility $x_t\sigma_z$ implies procyclical price per unit of variance risk $\gamma_{n,t}$
- Sharpe ratios decrease with maturity n because unpriced discount-rate risk $B_{x,n-1}^2\sigma_x^2$ rises

Model: Calibration and Simulations

Averages: 10,000 Simulated 200-Year Paths

	Sharpe ratio			Mean-var. ratio $\gamma_{n,t}$		
Maturity n	Very bad	Bad	Good	Very bad	Bad	Good
1 year	0.45	0.43	0.35	1.90	2.14	2.37
10 years	0.44	0.41	0.32	1.77	1.94	1.98
30 years	0.43	0.40	0.30	1.69	1.83	1.79

Very bad times: x_t in top decile. Bad & good times: x_t above & below its mean. Parameters from Lettau–Wachter (2007) & Gormsen (2021).

- Sharpe ratios countercyclical & decreasing in maturity
- Price per unit of variance risk $\gamma_{n,t}$ lower in bad times & at long maturities
- Conditional excess volatility inflates the denominator (risk) but not the numerator (risk premium)

$\Longrightarrow$ unpriced excess vol. in bad times helps make sense of multiple asset pricing facts

[risk & return over time and across horizons of dividend term structure]

Roadmap

1. Introduction
2. Theory: Priced & Unpriced Variance
3. Excess Volatility in Bad Times
4. Implications & A Simple Model
5. Final Notes

Final Notes

Summary:

- Conditional excess volatility: Market variance not compensated by a risk premium, inferred from the price per unit of variance risk without knowledge of fundamental value
- Procyclical price per unit of variance risk γ_t implies higher excess volatility in bad times
- If the true price of risk doubles in recessions, at least half of market variance is unpriced
- Implications for portfolio choice, volatility timing, real investment, & the equity term structure

Thank you!

Appendix

Priced Shocks Can Generate Excess Volatility

Example: Uncorrelated market & private-equity shocks X, Y, each with variance v; SDF $M = 1 - \gamma(X + Y)$

$$X = \underbrace{\frac{X+Y}{2}}_{s} + \underbrace{\frac{X-Y}{2}}_{u} \qquad \Longrightarrow \qquad \mu = \gamma v, \quad \psi = \frac{1}{2}, \quad \widetilde{\gamma} = 2\gamma$$

[spanned] [orthogonal]

- Both primitive shocks are priced, but half of market variance is orthogonal to the SDF
- In general, $\psi_t = 1 - \mathrm{Corr}_t^2(R_{m,t+1}, M_{t+1})$

Back to main

Asset Pricing Models: Setup

Priced dividend shocks: Equity as a levered claim to consumption, $D_t = C_t^\lambda$ [Campbell 1986; Abel 1999]

				Levered: ψ_t (%)		Published: ψ_t (%)	
	λ	Vol. (%)	Sharpe	Good	Bad	Good	Bad
Campbell–Cochrane	2.5	15.5	0.43	0.0	0.1	67	45
Bansal–Yaron	3	14.6	0.55	2.0	1.7	64	65
Drechsler–Yaron	2.5	17.2	0.52	5.3	11	73	47
Wachter	2.6	20.0	0.39	0.1	1.1	0.1	1.1
Lettau–Wachter		20.7	0.43			88	86

- λ from each paper's leverage parameter [Campbell–Cochrane: Set to match premium & volatility]
- All dividend shocks are priced, but dividend growth is too smooth [vol. 3–8%, data 11–14%]
- Published ψ_t mostly from unpriced shocks, except for Wachter
- Logs: $\gamma_t = -\mathrm{Cov}_t(\log M_{t+1}, \log R_{m,t+1})/\mathrm{Var}_t(\log R_{m,t+1})$, $\psi_t = 1 - \mathrm{Corr}_t^2(\log M_{t+1}, \log R_{m,t+1})$
- In levels: Realized γ_t nearly identical, but Wachter's $\psi_t = 75\%$ [SDF far more convex than returns]

Back to main

Asset Pricing Models: Issues with Volatility

Published calibrations: Annual volatilities vs. U.S. data [data: Beeler & Campbell 2012]

	Return vol. (%)	Log P/D vol.	Dividend growth vol. (%)
U.S. data, 1930–2008	20.2	0.45	11.1
Campbell–Cochrane	15.1	0.17	11.2
Bansal–Yaron	16.7	0.18	11.2
Drechsler–Yaron	17.4	0.17	11.0
Wachter (no-disaster years)	17.7	0.38	5.2
Lettau–Wachter	19.2	0.38	14.4

- U.S. data: Discount-rate news about twice as volatile as cash-flow news [Campbell 1991]
- Habit & long-run risks models: P/D too smooth $\Longrightarrow$ too little discount-rate news
- Return vol. close to data, but too much of it from unpriced dividend shocks
  [Except for Wachter (all shocks priced) & Lettau–Wachter (unpriced shocks to the price of risk)]

Back to main

Option-Implied Mean-Variance Ratio: Method

- Risk-neutral density each month from S&P 500 options, using a spline approach [Figlewski 2018]
- Projection of the SDF onto the market $\Longrightarrow$ physical density $\propto$ risk-neutral density $\times R_{m,t+1}^{\gamma_t}$
- Constant c in $\delta_t = 1/R_t^f + c$ & γ_t chosen so that realized returns are consistent with the implied physical densities, by likelihood-ratio test [Berkowitz 2001]
- Each physical density must integrate to one, which gives γ_t for a given δ_t

Back to main

Expected Mean-Variance Ratio: Machine-Learning Forecast

Risk premium forecast: From 100 anomaly portfolio returns, estimated via combination elastic net [Dong et al. 2022]

Expected γ_t Under Two Risk Premium Forecasts

	After market lows, $\gamma_B - \gamma_G$		γ_t on	
Risk premium forecast	1985–2017	1996–2017	NFCI	Recession prob.
Kelly–Pruitt	−2.12***	−2.60**	−0.96***	−1.99**
	(0.74)	(0.92)	(0.30)	(0.89)
Machine learning	0.50	−1.51	−1.31***	−1.56
	(1.56)	(1.11)	(0.29)	(1.01)

U.S., 1985–2017. HAR standard errors in parentheses.

- Better return forecast, but nearly uncorrelated with Kelly–Pruitt [R^2 2.8% vs. 1.0%; corr. 0.12]
- Main results: Same for financial conditions, same (or slightly weaker) for recessions

Back to main

Expected Mean-Variance Ratio: Alternative Variance Forecasts

Variance forecast: HAR-RV [Corsi 2009] & EGARCH [Nelson 1991], estimated in real time from daily returns

Expected γ_t Under Three Variance Forecasts

	After market lows, $\gamma_B - \gamma_G$		γ_t on	
	Post-1964	Post-1996	NFCI	Recession prob.
AR(1)	-1.77^{***}	-2.23^{***}	-0.65^{***}	-2.18^{***}
	(0.41)	(0.66)	(0.20)	(0.69)
HAR-RV	-1.72^{***}	-2.17^{***}	-0.65^{***}	-2.14^{***}
	(0.39)	(0.61)	(0.19)	(0.66)
EGARCH	-2.71^{***}	-2.90^{***}	-0.98^{***}	-3.37^{***}
	(0.56)	(0.69)	(0.27)	(0.77)

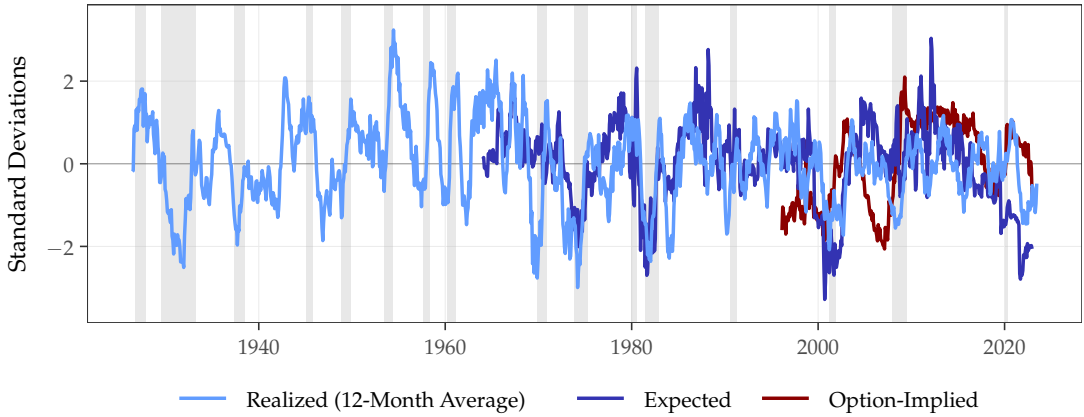
Standard errors are HAR [Lazarus et al. (2018)].

- Both more accurate than the AR(1) [squared error 5% & 10% lower]
- Main results: Same with HAR-RV, stronger with EGARCH

Back to main

Estimates of Sharpe Ratios

Standardized U.S. Sharpe ratios:



Back to main

Market Volatility in Recent Crises

With a constant price of risk: $\mu_t = \widetilde{\gamma}\,\sigma_t^2$

	Volatility (%)			Implied μ_t (%)	
	12 months before	Peak	Variance ratio	$\widetilde{\gamma} = 1$	$\widetilde{\gamma} = 3$
Covid 2020	15.4	99.3	41.5	99	296
Financial crisis 2007–2009	14.9	83.2	31.2	69	207
Dot-com 2001	22.9	36.9	2.6	14	41

Back to main

Consumption-Based Evidence

Power utility, risk aversion γ_c: $\gamma_t = -\mathrm{Cov}_t(\log M_{t+1}, \log R_{m,t+1})/\sigma^2_{m,t} = \gamma_c\, \rho_{m,c,t}\, \sigma_{c,t}/\sigma_{m,t}$

U.S. Quarterly Estimates

	Normal times	NBER recessions
Correlation between consumption growth & market return, $\rho_{m,c}$	0.18	0.14
Consumption volatility, σ_c	0.42%	0.55%
Market volatility, σ_m	6.6%	11.3%
γ_t, power utility ($\gamma_c = 100$)	1.12	0.69
γ_t, recursive utility ($\gamma_c = 10$, IES $= 4/3$)	6.79	5.40

- Consumption volatility 30% higher in recessions, market volatility 70% higher
- Procyclical price per unit of variance risk γ_t under both utility specifications

Back to main

Investment Regressions

Regressions of real investment growth on indicator for low expected price per unit of variance risk γ_q:

$$\Delta \text{Inv}_{q,q+s} = \alpha_s + \beta_s \mathbb{1}\{\gamma_q^{\text{expected}} < 25\text{th pctile}\} + \text{controls}_q + \varepsilon_{q,q+s}$$

Coefficient β_s: Real Private Investment Growth

Horizon (quarters)	$s=1$	$s=2$	$s=3$	$s=4$	$s=5$	$s=6$	$s=7$	$s=8$
No controls	-0.01*	-0.02*	-0.03	-0.03	-0.04	-0.03	-0.04	-0.04
	(0.01)	(0.01)	(0.02)	(0.03)	(0.03)	(0.04)	(0.04)	(0.04)
With controls	-0.02***	-0.02	-0.02	-0.03	-0.03	-0.03	-0.04	-0.03
	(0.01)	(0.01)	(0.03)	(0.03)	(0.04)	(0.04)	(0.04)	(0.04)

Controls: VIX, recession probability, NFCI, dividend-price ratio, credit spread. HAR standard errors in parentheses.

Back to main