On the Moments of the Stochastic Discount Factor

Ian Martin   Ran Shi

March 2026

Background

- Asset prices are often used for assessing expectations:
  - forward rates $\longrightarrow$ risk-neutral expected future interest rates
  - breakeven inflation $\longrightarrow$ risk-neutral expected future inflation
  - CDS rates $\longrightarrow$ risk-neutral default probabilities
  - implied volatility $\longrightarrow$ risk-neutral volatility
  - ...
- These risk-neutral quantities are almost continuously observable
- They are model-free
- And they embody the collective views of market participants
- But they are distorted by risk: people pay more for assets that pay off in bad states

Background (2)

- Hansen and Jagannathan (1991) introduced the idea that the importance of risk considerations can be captured in a general way via the SDF
- If SDF is constant, risk-neutral and true distributions are the same thing
- If SDF is volatile, risk matters a lot
- HJ bound: Sharpe ratios provide lower bounds on SDF volatility
- “Important direction” for future research: move beyond mean and variance of the SDF

Background (3)

- Snow (1991): lower bounds on the θth moment of the SDF in terms of the $\frac{\theta}{\theta-1}$th moment of returns
- So when θ is close to 1, the Snow bound depends on extremely high return moments, which are hard to estimate empirically
- This might seem surprising: if a riskless asset is traded, we can perfectly infer the first moment of the SDF!
- Why is it so hard to restrict nearby moments?

This paper

- New bounds on arbitrary moments of the SDF
- Like the prior literature, we exploit the true distribution of returns, which we infer from the time series
- Unlike the prior literature, we also exploit the risk-neutral distribution, which is observable from option prices
- This allows us to derive stronger bounds than the prior literature, and they are useful even when θ is close to 1

The bad news

- We find that the moments of the SDF grow very rapidly as θ rises above one
- A singularity occurs at around $\theta = 1.7$
- Our results suggest that the SDF may have infinite volatility
- Bad news for the vast literature based on mean–variance analysis!
  - Markowitz (1952), Sharpe (1964), Lintner (1965), Black, Jensen and Scholes (1972), Fama and MacBeth (1973), Chamberlain and Rothschild (1983), Hansen and Richard (1987), Gibbons, Ross and Shanken (1989), Fan, Liao and Yao (2015), Kozak, Nagel and Santosh (2020), Bryzgalova, Huang and Julliard (2023), Chernov, Kelly, Malamud and Schwab (2025), . . .

The good news

- But we also show that variance bounds are inherently unstable as a matter of theory
- As are θth moment bounds if $\theta < 0$ or $\theta > 1$
- But intermediate bounds, $\theta \in (0, 1)$, and entropy bounds are well-behaved
- And they have nice economic interpretations
  - Measure of attractiveness of investment opportunities (analogous to HJ link between SDF variance and Sharpe ratios)
  - Measure of market risk aversion (generalizing classic results of Merton (1969) and Samuelson (1969))
- These well-behaved measures take plausible values in the data

The cumulant-generating function

- Fix a return R_{t+1}, riskless rate $R_{f,t+1}$, and SDF M_{t+1}, and define

$$\boldsymbol{\kappa}_t(\theta_1, \theta_2) \equiv \log \mathbb{E}_t \left[\left(M_{t+1} R_{f,t+1} \right)^{\theta_1} \left(R_{t+1} / R_{f,t+1} \right)^{\theta_2} \right]$$

- Notice that

$$\boldsymbol{\kappa}_t(0,0) = 0, \qquad \boldsymbol{\kappa}_t(1,0) = 0, \qquad \boldsymbol{\kappa}_t(1,1) = 0$$

- $\boldsymbol{\kappa}_t(0, \theta_2) = \log \mathbb{E}_t \left[\left(R_{t+1} / R_{f,t+1} \right)^{\theta_2} \right]$ — true return cumulants
- $\boldsymbol{\kappa}_t(1, \theta_2) = \log \mathbb{E}^*_t \left[\left(R_{t+1} / R_{f,t+1} \right)^{\theta_2} \right]$ — risk-neutral cumulants
- $\boldsymbol{\kappa}_t(\theta_1, 0) = \log \mathbb{E}_t \left[\left(M_{t+1} R_{f,t+1} \right)^{\theta_1} \right]$ — SDF cumulants

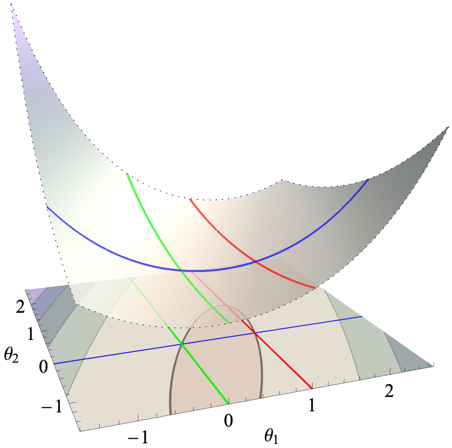

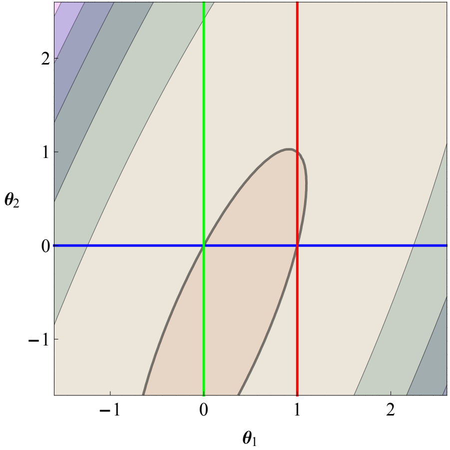

- Why bother to take logs? Because the CGF is a convex function of θ_1 and θ_2
- (The MGF is too, but this is a much cruder—and therefore less useful—property)

Observable slices of the CGF

- We estimate the true distribution from the time series of realized returns
- The risk-neutral distribution is revealed by option prices:

$$\boldsymbol{\kappa}_t\,(1,\theta) = \log \mathbb{E}_t^* \left[\left(R_{t+1}/R_{f,t+1} \right)^{\theta} \right]$$

$$= \log \left\{ 1 + \theta(\theta-1) \left[\underbrace{\int_0^1 K^{\theta-2}\, \mathrm{put}_t(K R_{f,t+1})\, \mathrm{d}K + \int_1^\infty K^{\theta-2}\, \mathrm{call}_t(K R_{f,t+1})\, \mathrm{d}K}_{\text{price of a portfolio of options with different strikes } K} \right] \right\}$$

- Empirically, the risk-neutral distribution is easier to estimate than the true distribution

Example 1: joint lognormality

- $M_{t+1}R_{f,t+1} = e^{-\frac{1}{2}\lambda^2 - \lambda Z}$, $R_{t+1}/R_{f,t+1} = e^{\mu - \frac{1}{2}\sigma^2 + \sigma W}$
- The CGF is quadratic:

$$\boldsymbol{\kappa}_t(\theta_1, \theta_2) = \mu\theta_2(1-\theta_1) + \frac{1}{2}\lambda^2\theta_1(\theta_1 - 1) + \frac{1}{2}\sigma^2\theta_2(\theta_2 - 1)$$

- Observable slices:

$$\boldsymbol{\kappa}_t(0, \theta_2) = \mu\theta_2 + \frac{1}{2}\sigma^2\theta_2(\theta_2 - 1)$$
$$\boldsymbol{\kappa}_t(1, \theta_2) = \frac{1}{2}\sigma^2\theta_2(\theta_2 - 1)$$

- Almost all well-known equilibrium models look like this
- Risk-neutral CGF (calculated from option prices) is uninteresting

Example 2: adding jumps

- Add Poisson jumps arriving at rate ω, SDF jump size J_1, return jump size J_2
- Observable slices:

$$\boldsymbol{\kappa}_t(0, \theta_2) = \cdots + \omega \left[(1 + J_2)^{\theta_2} - \theta_2 J_2 - 1 \right]$$
$$\boldsymbol{\kappa}_t(1, \theta_2) = \cdots + \omega(1 + J_1) \left[(1 + J_2)^{\theta_2} - \theta_2 J_2 - 1 \right]$$

- Now options are useful: risk-neutral distribution reveals J_1 (the SDF jump size) which is not revealed by the true distribution
- All moments are finite

Example 3: learning about model parameters

- Geweke (2001) and Weitzman (2007) have shown that things change dramatically when agents must learn model parameters
- If agents are uncertain about ω, with prior $\omega \sim \text{Exp}(\overline{\omega})$, and jumps are bad news ($J_1 > 0, J_2 < 0$), then true and risk-neutral return moments diverge for sufficiently negative θ
- And the θth moment of SDF diverges once θ exceeds a critical value > 1

Example 4: heterogeneous beliefs

- Brownian limit of Martin and Papadimitriou (2022)
- Observable slices:

$$\boldsymbol{\kappa}_t(0,\theta_2) = \frac{1+\delta}{2\delta}\sigma^2\theta_2 + \frac{1}{2}\sigma^2\theta_2^2$$
$$\boldsymbol{\kappa}_t(1,\theta_2) = \frac{1+\delta}{2\delta}\sigma^2\theta_2(\theta_2-1)$$

- SDF moments:

$$\boldsymbol{\kappa}_t(\theta,0) = \frac{(1+\delta)^2\theta(\theta-1)\sigma^2}{2\delta(1+\delta-\theta)} + \frac{1}{2}\log\frac{1+\delta}{1+\delta-\theta} - \frac{1}{2}\theta\log\frac{1+\delta}{\delta}$$

- $(1+\delta)$th and higher SDF moments are unbounded
- With sufficient heterogeneity, $\delta \leq 1$, SDF variance is infinite: arbitrarily high Sharpe ratios are attainable via option strategies

Observable slices of the CGF

- Equilibrium models pin down the entire CGF surface
- Empiricists don't have this luxury!
- What can we say given knowledge only of the true distribution $\boldsymbol{\kappa}_t(0, \theta_2)$ and risk-neutral distribution $\boldsymbol{\kappa}_t(1, \theta_2)$?

Result (New moment bounds)

For $\theta < 0$ or $\theta > 1$:

$$\mathbb{E}_t\left[\left(M_{t+1}R_{f,t+1}\right)^{\theta}\right] \geq \sup_y \left\{\mathbb{E}_t^*\left[\left(\frac{R_{t+1}}{R_{f,t+1}}\right)^{y}\right]\right\}^{\theta}\left\{\mathbb{E}_t\left[\left(\frac{R_{t+1}}{R_{f,t+1}}\right)^{\frac{\theta}{\theta-1}y}\right]\right\}^{1-\theta}$$

For $\theta \in (0,1)$:

$$\mathbb{E}_t\left[\left(M_{t+1}R_{f,t+1}\right)^{\theta}\right] \leq \inf_y \left\{\mathbb{E}_t^*\left[\left(\frac{R_{t+1}}{R_{f,t+1}}\right)^{y}\right]\right\}^{\theta}\left\{\mathbb{E}_t\left[\left(\frac{R_{t+1}}{R_{f,t+1}}\right)^{\frac{\theta}{\theta-1}y}\right]\right\}^{1-\theta}$$

- $y = 1$ recovers Snow (and fails to exploit information in the risk-neutral distribution)
- Bounds hold conditionally and unconditionally

Result (New moment bounds)

For $\theta < 0$ or $\theta > 1$:

$$\kappa_t(\theta, 0) \geq \sup_y \; \theta\, \kappa_t(1, y) + (1-\theta)\, \kappa_t\left(0, \frac{\theta}{\theta - 1} y\right)$$

For $\theta \in (0,1)$:

$$\kappa_t(\theta, 0) \leq \inf_y \; \theta\, \kappa_t(1, y) + (1-\theta)\, \kappa_t\left(0, \frac{\theta}{\theta - 1} y\right)$$

- $y = 1$ recovers Snow (and fails to exploit information in the risk-neutral distribution)
- Bounds hold conditionally and unconditionally

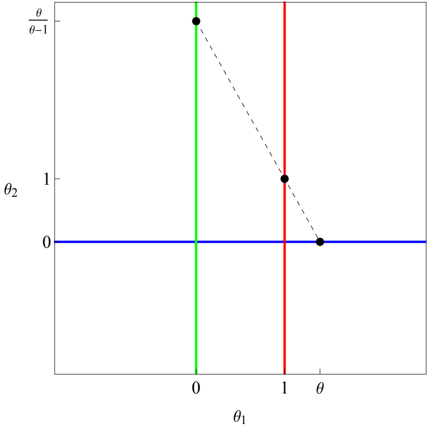

$y = 1$: Snow, looking through $(1, 1)$

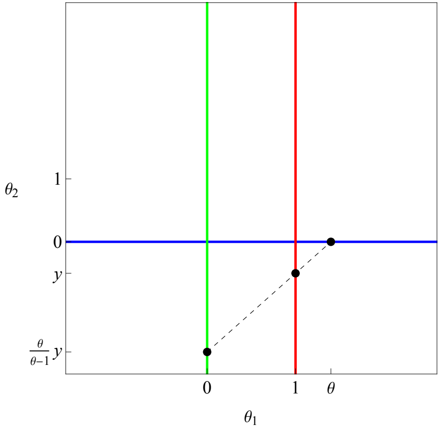

General y

- **Proof idea**: Exploit convexity via three collinear points $(\theta, 0)$, $(0, \frac{\theta}{\theta-1}y)$, $(1, y)$
- Conventional approach implicitly fixes $y = 1$. Free parameter y lets us "look around"

Summary so far

- We have derived new bounds on the moments of the SDF that play the true and risk-neutral distributions off against one another
- Couldn't we just do this by plugging returns on option strategies into, say, the Hansen–Jagannathan bound (Liu, 2021)?
  - We only observe option prices over a short time series
  - Option strategies are highly skewed and fat-tailed, so very sensitive to outliers
- Our approach allows us to sidestep this issue
- We use option prices, not option returns

Conditional and unconditional CGFs

- These bounds apply conditionally or unconditionally
- Conditional risk-neutral CGF $\boldsymbol{\kappa}_t(1,\theta)$ is easy to measure

$$\boldsymbol{\kappa}_t\left(1,\theta\right) = \log \mathbb{E}_t^* \left[\left(R_{t+1}/R_{f,t+1} \right)^{\theta} \right]$$

$$= \log \left\{ 1 + \theta(\theta-1) \left[\underbrace{\int_0^1 K^{\theta-2}\, \text{put}_t(KR_{f,t+1})\, \mathrm{d}K + \int_1^\infty K^{\theta-2}\, \text{call}_t(KR_{f,t+1})\, \mathrm{d}K}_{\text{price of a portfolio of options with different strikes } K} \right] \right\}$$

- But conditional *true* distribution is hard to measure, so we work unconditionally
- As we measure conditional risk-neutral distribution perfectly, most statistical uncertainty is associated with the true distribution, not the risk-neutral distribution

Conditional and unconditional CGFs

- Unconditional true CGF of returns is

$$\boldsymbol{\kappa}(0,\theta) = \log \mathbb{E}\left[\left(R_{t+1}/R_{f,t+1}\right)^{\theta}\right]$$

- Assuming stationarity and ergodicity, we use the time-series estimator

$$\boldsymbol{\kappa}(0,\theta) = \log \frac{1}{T}\sum_{t=0}^{T-1}\left(R_{t+1}/R_{f,t+1}\right)^{\theta}$$

Conditional and unconditional CGFs

- Unconditional risk-neutral CGF is

$$\begin{aligned}\boldsymbol{\kappa}(1,\theta) &= \log \mathbb{E}\left[M_{t+1}R_{f,t+1}\left(R_{t+1}/R_{f,t+1}\right)^{\theta}\right]\\ &= \log \mathbb{E}\left(\mathbb{E}_t\left[M_{t+1}R_{f,t+1}\left(R_{t+1}/R_{f,t+1}\right)^{\theta}\right]\right)\\ &= \log \mathbb{E}\left(\underbrace{\mathbb{E}_t^*\left[\left(R_{t+1}/R_{f,t+1}\right)^{\theta}\right]}_{\exp\{\boldsymbol{\kappa}_t(1,\theta)\}}\right)\end{aligned}$$

- Again, we use the standard time-series estimator, replacing $\mathbb{E}$ with $\frac{1}{T}\sum_t$

$$\boldsymbol{\kappa}(1,\theta) = \log \frac{1}{T}\sum_{t=1}^{T} \exp\left\{\boldsymbol{\kappa}_t(1,\theta)\right\}$$

Finite samples

- You might think that the relatively short sample of option prices is a problem
- But actually, even after 150 years, most of the estimation uncertainty is due to the realized return series
- Risk-neutral prob. of 10% market decline in one month:
  $\widehat{\Pr}^* = 6.56\%$, s.e. $= 0.22\%$ (26 years of options)
- True probability of 10% decline in one month:
  $\widehat{\Pr} = 1.82\%$, s.e. $= 0.32\%$ (150 years of returns!)
- We use option prices to infer the risk-neutral distribution
- We do not use option returns, which are extremely noisy

Data

- Market returns:
  - Global Financial Data: 1871–1926
  - CRSP: 1926–2022

  Baseline sample is monthly from 1872 to 2022
- Jordà–Schularick–Taylor Macrohistory Database (1872–2020): annual returns from Jordà, Knoll, Kuvshinov, Schularick, and Taylor (2019)
- S&P 500 index options from OptionMetrics
- We always use whichever of bid and offer prices give the weaker bounds. Mid-market prices would make our results even more dramatic
- We extrapolate outside the range of observed strikes via no arbitrage arguments alone
- Robust to: extrapolating with flat vol smile; not extrapolating at all; using daily rather than monthly returns; using only options post-2008; alternative moneyness filters

The two observable slices of the CGF in the data

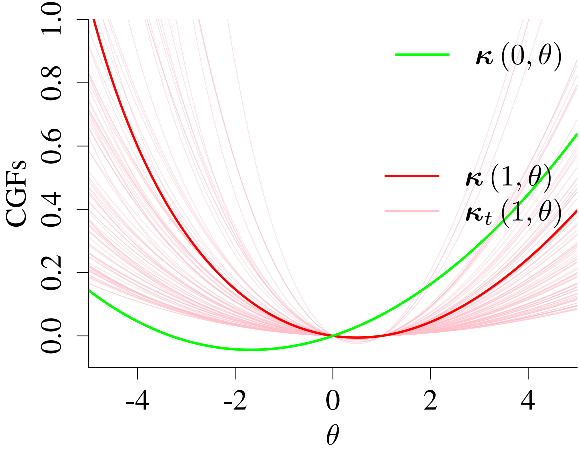

The convexity bounds for the SDF moments

Realized returns: 1872–2022

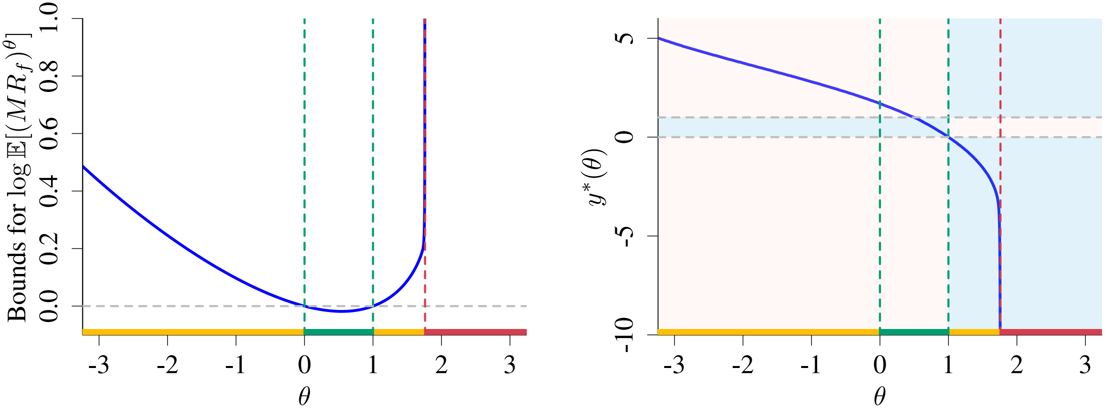

The convexity bounds for the SDF moments

Realized returns: 1946–2022

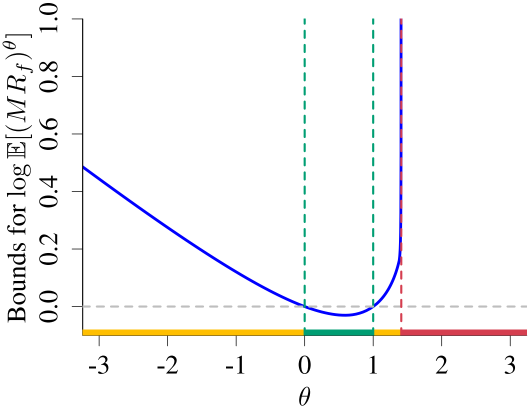

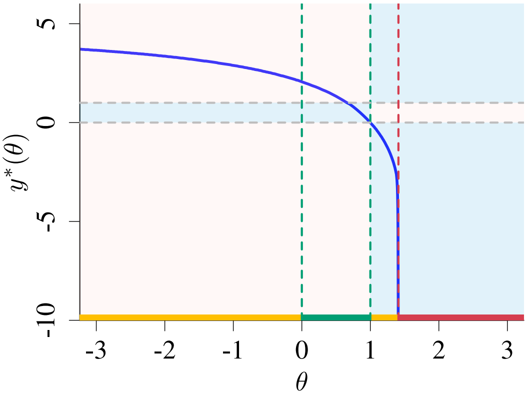

The singularity is near

Over the one-month horizon

sample	est.	bootstrap CI	$\mathbb{E}$-bootstrap CI	$\mathbb{E}^*$-bootstrap CI
1872-2022	1.72	(1.52, 2.05)	(1.50, 1.97)	(1.62, 1.87)
1946-2022	1.38	(1.20, 1.60)	(1.20, 1.56)	(1.34, 1.45)
1996-2022	1.44	(1.23, 1.78)	(1.23, 1.77)	(1.39, 1.52)

- The singularity emerges below two
- The majority of estimation uncertainty comes from the realized returns

The singularity is near

Over the one-year horizon

sample	est.	bootstrap CI	$\mathbb{E}$-bootstrap CI	$\mathbb{E}^*$-bootstrap CI
1872-2022	1.67	$(1.27, 2.50)$	$(1.27, 2.05)$	$(1.60, 2.10)$
1946-2022	1.28	$(1.15, 1.50)$	$(1.14, 1.42)$	$(1.23, 1.38)$
1996-2022	1.36	$(1.14, 1.92)$	$(1.13, 1.73)$	$(1.32, 1.52)$
JKKST annual	1.36	$(1.23, 1.56)$	$(1.21, 1.52)$	$(1.32, 1.51)$

Interpretation

- Variance bounds are inherently prone to explosion even in population: the lower bound for $\log \mathbb{E}\left[(MR_f)^2\right]$ is

$$\sup_{y} \; 2\,\boldsymbol{\kappa}(1, y) - \boldsymbol{\kappa}(1, 2y)$$

- Requires maximizing the difference between two convex functions—badly behaved!
- Also true for any $\theta > 1$ or $\theta < 0$
- For example, in the MP model, with $\theta = 2$,

$$\boldsymbol{\kappa}_t(2, 0) \geq \frac{1-\delta}{\delta}\sigma^2 y^2 - \frac{2(1+\delta)}{\delta}\sigma^2 y$$

  - If $\delta \leq 1$, bound diverges as $y \to -\infty$. This accurately reflects arbitrarily high Sharpe ratios (attained by short option positions in the model, consistent with our findings)

Finite sample issues when $\theta \notin (0,1)$

- If there are puts (with positive bid price) in the dataset whose strikes are lower than the lowest observed return in sample, then the lower bound can be made arbitrarily large by sending y to $\pm\infty$
- The worst monthly return is -29% (September 1931)
- Put options with strikes more than 29% out of the money have positive *bid* prices on 73% of days in our sample
- The only robust response: focus on $\theta \in (0,1)$
- We want to connect to prior literature, so seek interior optima in y when $\theta \notin (0,1)$
- But even this approach breaks down above the singularity: the bounds increase monotonically with y

In contrast, intermediate moments are well-behaved

- For $\theta \in (0, 1)$, moment bounds are optimized by minimizing a weighted average of two convex functions
- This is a well-behaved procedure
- Unfamiliar but reliable
  - Upper bound on $\mathbb{E}\sqrt{M_{t+1}R_{f,t+1}} \Rightarrow$ lower bound on var $\sqrt{M_{t+1}R_{f,t+1}}$
  - Natural interpretation via willingness-to-pay

And intermediate bounds have nice economic interpretation

- The HJ bound connects SDF variance to an easily interpretable measure of the attractiveness of investment opportunities: the Sharpe ratio
- The intermediate bounds have a similar property
- Take the perspective of a one-period CRRA investor, risk aversion γ
- The attractiveness of investment opportunities can be quantified using the willingness-to-pay (WTP)—the fraction of initial wealth the investor would sacrifice to be allowed to trade risky assets, g_γ
- Economically, this makes more sense than the Sharpe ratio, because it does not require us to adopt the (extremely unreasonable!) perspective of a mean–variance investor

Willingness-to-pay to trade risky assets

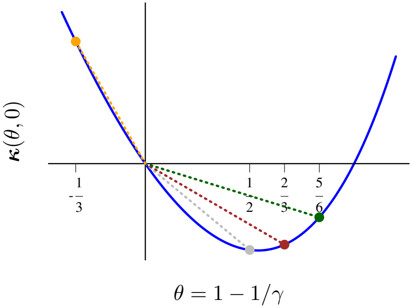

Moment bounds

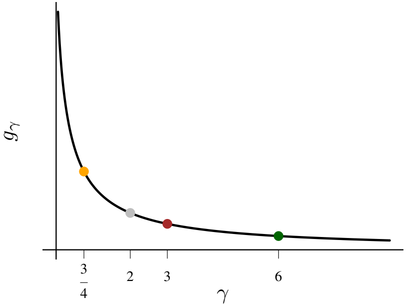

WTP against risk aversion γ

- Moment bounds $B(\theta)$ translate into bounds on WTP: $g_\gamma \geq \left| \frac{B(1-1/\gamma)}{1-1/\gamma} \right|$

$\theta \in (0,1)$: moment bounds (1872–2022)

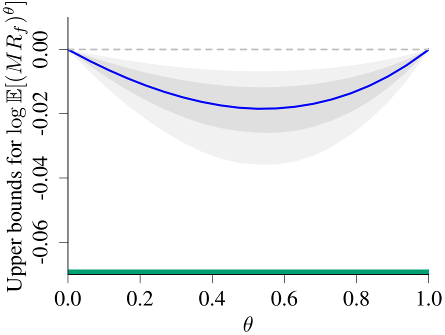

Moment bounds

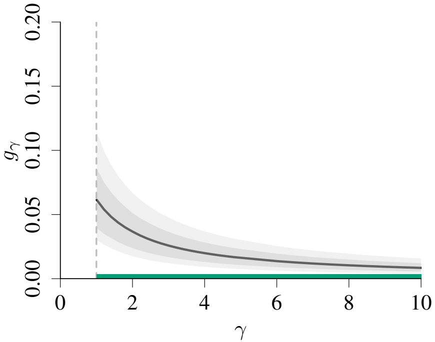

WTP against risk aversion γ

- Grey bands: 68% and 95% confidence intervals

$\theta \in (0, 1)$: moment bounds (1946–2022)

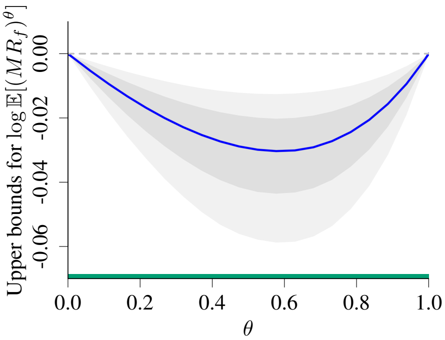

Moment bounds

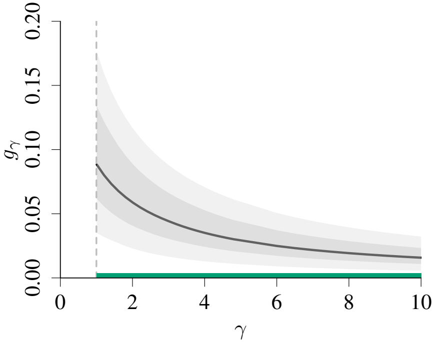

WTP by risk aversion γ

- Log investor ($\gamma = 1$): would sacrifice $\geq 9\%$ of wealth
- $\gamma = 2$: $\geq 6\%$

Entropy measures

- At the extremes of the well-behaved interval, $\theta \in (0,1)$, the bounds are trivial: as $\mathbb{E}\left(M_{t+1}R_{f,t+1}\right) = 1$ and $\mathbb{E}\left(M_{t+1}R_{t+1}\right) = 1$, we have $\boldsymbol{\kappa}(1,0) = \boldsymbol{\kappa}(1,1) = 0$
- But the slopes at each end give bounds on Theil's two entropy measures

$$L^{(1)}(M_{t+1}R_{f,t+1}) = \mathbb{E}^* \log(M_{t+1}R_{f,t+1}) \quad \text{and} \quad L^{(2)}(M_{t+1}R_{f,t+1}) = -\,\mathbb{E}\log(M_{t+1}R_{f,t+1})$$

- First measure can be written as $\mathbb{E}^* \log \frac{d\mathbb{Q}}{d\mathbb{P}}$ (used by Hansen and Sargent as a penalty on belief distortions; cf also Chen, Hansen and Hansen, 2020)
  - Of all the well-behaved moment and entropy measures, this is the most sensitive to the right tail of the SDF
- Second measure is the focus of Backus, Chernov and Zin (2014) in rep agent models

Result (Entropy bounds)

$$L^{(1)}\left(M_{t+1}R_{f,t+1}\right) \geq \sup_y \; y\,\mathbb{E}^*\log\left(R_{t+1}/R_{f,t+1}\right) - \log\mathbb{E}\left[\left(R_{t+1}/R_{f,t+1}\right)^y\right]$$

$$L^{(2)}\left(M_{t+1}R_{f,t+1}\right) \geq \sup_y \; y\,\mathbb{E}\log\left(R_{t+1}/R_{f,t+1}\right) - \log\mathbb{E}^*\left[\left(R_{t+1}/R_{f,t+1}\right)^y\right]$$

- These bounds are dual to one another: they interchange the role of true and risk-neutral distributions
- They are both well-behaved (sup over linear function of y minus convex function of y)
- $L^{(1)}$ bound is entirely new
- $L^{(2)}$ case specializes to Bansal–Lehmann bound if we fix $y = 1$ to avoid using options
- The optimizing values of y in these bounds have a nice interpretation

The gap between the two entropy measures

$$L^{(1)}\left(M_{t+1}R_{f,t+1}\right) - L^{(2)}\left(M_{t+1}R_{f,t+1}\right) = \frac{\kappa_3}{6} + \frac{\kappa_4}{12} + \frac{\kappa_5}{40} + \cdots$$

- Difference between the two measures is related to higher cumulants of log SDF: skewness, excess kurtosis ...
- Under lognormality: $L^{(1)} = L^{(2)} = \frac{1}{2}\operatorname{var}\log\left(M_{t+1}R_{f,t+1}\right)$
- A positive gap reveals right-skewness and fat tails of log SDF

Merton–Samuelson, generalized

- The optimizing values of y in the entropy bounds have a nice interpretation
- Consider a myopic power utility investor holding the S&P 500
- If R/R_f is lognormal, the Merton–Samuelson (1969) calculation yields:

$$\gamma = \frac{\mu - r_f}{\sigma^2} \approx 2$$

- Under our framework, without any distributional assumptions:
  - y^* optimizing $L^{(1)}$ bound is $-\gamma$
  - y^* optimizing $L^{(2)}$ bound is γ

First entropy metric $L^{(1)}$

sample	est.	boot CI	$\mathbb{E}$-boot	$\mathbb{E}^*$-boot
1872 – 2022	0.088	(0.033, 0.175)	(0.033, 0.177)	(0.084, 0.093)
1946 – 2022	0.173	(0.070, 0.344)	(0.071, 0.348)	(0.165, 0.182)
1996 – 2022	0.157	(0.016, 0.445)	(0.016, 0.439)	(0.149, 0.165)

- Stable across subsamples
- Most uncertainty is associated with the time series of realized returns

Second entropy metric $L^{(2)}$

					B–L $L^{(2)}\|_{y=1}$	
sample	est.	boot CI	$\mathbb{E}$-boot	$\mathbb{E}^*$-boot	est.	boot CI
1872 – 2022	0.062	(0.023, 0.122)	(0.023, 0.120)	(0.060, 0.065)	0.052	(0.023, 0.080)
1946 – 2022	0.089	(0.038, 0.174)	(0.038, 0.174)	(0.084, 0.094)	0.066	(0.036, 0.100)
1996 – 2022	0.091	(0.009, 0.262)	(0.009, 0.258)	(0.087, 0.097)	0.067	(0.006, 0.125)

- Full sample: 6.2% (optimized) vs 5.2% (Bansal–Lehmann)
- Option prices raise the lower bound by $\sim$1 percentage point
- $L^{(1)} > L^{(2)}$ in all samples
- The log SDF is right-skewed

Implied risk aversion

	$L^{(1)}$		$L^{(2)}$	
horizon	est.	boot CI	est.	boot CI
1	2.36	(1.42, 3.38)	1.69	(1.00, 2.48)
2	2.21	(1.41, 3.11)	1.67	(1.05, 2.40)
3	2.14	(1.38, 3.08)	1.65	(1.02, 2.36)
4	2.17	(1.43, 3.10)	1.63	(1.09, 2.27)
5	2.15	(1.40, 3.14)	1.62	(1.03, 2.35)
6	2.07	(1.34, 3.12)	1.63	(1.03, 2.39)
9	1.84	(1.06, 2.95)	1.66	(1.00, 2.50)
12	1.74	(0.96, 2.99)	1.66	(0.99, 2.65)

Implications for mean–variance analysis

- If the SDF has infinite variance:
  - No well-defined maximum Sharpe ratio
  - Mean–variance framework collapses
  - Testing factor model efficiency breaks down
- Affects GRS, high-dimensional variants, ML-based SDF identification
- Recent papers use regularization to avoid spuriously high Sharpe ratios due to in-sample overfitting
- But our results suggest that high Sharpe ratios may be a feature of the population, not just an in-sample artifact
- It is problematic simply to ignore option prices. As Black and Scholes (1973) observed, "almost all corporate liabilities can be viewed as combinations of options"

Summary

Bad news:

- SDF moments diverge rapidly above $\theta = 1$; singularity < 2
- Variance bounds inherently unstable—even in population
- Problematic for mean–variance analysis!

Good news:

- Intermediate moments and entropy bounds well-behaved
- These are natural targets for machine learning approaches
- Natural interpretations: WTP, risk aversion
- Stable empirically; plausible values

Conclusions

1. Use entropy and $\theta \in (0, 1)$ moment bounds—not variance bounds
2. Option prices contain information about the SDF that returns do not reveal
3. Extreme Sharpe ratios may be a population feature, not overfitting
4. The empirical evidence makes perfect sense from the perspective of a CRRA investor