THE EQUITY RISK PREMIUM

AN OPTIONS-BASED APPROACH

ALAN L. LEWIS

BAYES BUSINESS SCHOOL, FEB. 19, 2025

Topics

- **Define the forward-looking Equity Risk Premium (ERP).**
- **My modeling approach: exponential tilt of broad-based US market return probability distributions (P/Q) by a parameter γ. (P = real-world; Q = risk-neutral). I present two rationales for the tilt approach.**
- **I show selected results from three option-based sub-projects: (1) a Covid-era study, (2) an 11-year study, (3) a third-Friday study.**

Topics (continued)

- **Estimation methods for the tilt parameter [= Coefficient of Relative Risk Aversion (CRRA) γ]: (1) Equity-based; (2) Options-based.**
- **Results from the Covid study: I show pandemic-era ERP term structures with $\gamma = 3 \pm 0.5$:**
- **Estimation methods for both ERP's and Q-distributions.**
- **Rationale #2 (for the tilt method) poses the empirical question: what is the best systematic forecast of the real-world (one-time event) probability density $P_{t,T}(x)$ relative to plausible alternatives? This is really <u>the key issue</u>, addressed in my Third-Friday study.**
- **Option data and automation notes.**
- **Conclusions.**

Defining the (US market) forward-looking Equity Risk Premium (ERP)

- $\mathrm{ERP}_{t,T} = \frac{100}{T-t}\left\{\mathrm{E}_t^P\left[\frac{\hat{S}_T}{\hat{S}_t}\right] - (1 + R_{t,T}^f)\right\}$,
- where $\hat{S}_t$ is a broad-based US total return equity index ("hats": dividends reinvested)
- $R_{t,T}^f$ is an observable risk-free return (US Treasury)

- Notes:
- As with interest rates, I always <u>annualize</u> and express results in percent.
- Expectations $\mathrm{E}_t^P[..]$ use a (best candidate) <u>Real-world</u> probability (aka P-probability, "<u>quasi</u>-objective", statistical, actuarial), conditioned on time-t information.
- $P_{t,T}(x)$ is a <u>one-time event</u> probability.

My modeling approach: exponential tilting

P/Q duality under exponential tilting/Esscher transform:

$$q_{X_T}(x) = \frac{e^{-\gamma x} p_{X_T}(x)}{\int e^{-\gamma x} p_{X_T}(x)\, dx} \quad \Leftrightarrow \quad p_{X_T}(x) = \frac{e^{\gamma x} q_{X_T}(x)}{\int e^{\gamma x} q_{X_T}(x)\, dx},$$

where $X_T = \log \frac{\hat{S}_T}{\hat{S}_t}$ = log-total-return, γ = tilt parameter, (t-dependence, too!)

P = real-world distribution with density p(x),

Q = risk-neutral distribution (inferred from option prices).

Note: same tilt form works for log-price-return densities

Rationale #1: Representative Agent Argument
γ = Coefficient of Relative Risk Aversion (CRRA)

Suppose an Investor Representative Agent with power utility of wealth: $U(W_T) = \frac{e^{-\rho T}}{1-\gamma}\left(\frac{W_T}{W_0}\right)^{1-\gamma}$ **...**

... Chooses a portfolio from a (total return) equity index with price $\hat{S}_0$, risk-free discount bond contract with price $D_{0,T}$, and an equity derivative with generic payoff $V_T = w(x_T)$, where $x_T = \log \hat{S}_T$.

Agent's objective: maximize P-expected utility (real-world expected utility).

Market clearing equilibrium: 100% equity allocation, zero to the rest.

Result: $V_0 = D_{0,T}\, E_0^Q[V_T]$,

where $E_0^Q[V_T] = \frac{E_0^P[m_{0,T} V_T]}{E_0^P[m_{0,T}]} = \frac{E_0^P\left[\left(\frac{\hat{S}_T}{\hat{S}_0}\right)^{-\gamma} V_T\right]}{E_0^P\left[\left(\frac{\hat{S}_T}{\hat{S}_0}\right)^{-\gamma}\right]}$, **($m_{0,T}$ is the "stochastic discount factor")**

Finally, specialize to: $w(x) = \delta(x - y - \log \hat{S}_0)$ **(Dirac delta)** $\Rightarrow$

$$q_{X_T}(x) = \frac{e^{-\gamma x}\, p_{X_T}(x)}{\int e^{-\gamma x} p_{X_T}(x)\, dx},$$ **where recall** $X_T = \log \frac{\hat{S}_T}{\hat{S}_0}$.

Rationale #2: Some Q to P transformation exists by the FTAP

Premises: (i) a real-world $P_{t,T}(x)$ exists; (ii) There are no arbitrage opportunities.

Then, (by the Fundamental Theorem of Asset Pricing) securities are priced (relative to a risk-free bond) by <u>some</u> equivalent martingale-measure $Q_{t,T}(x)$. So, under the premises a $Q \leftrightarrow P$ transformation exists. (Here x's are broad-index log-price-returns).

We can (essentially) observe $Q_{t,T}(x)$ via option market quotes

$P_{t,T}^{(\gamma)}(x) = C\, e^{\gamma x}\ Q_{t,T}(x)$ are a reasonable class of one-parameter candidates for $P_{t,T}$ that deserve to be tested for their forecasting abilities against plausible alternatives (Null Hypotheses H_0). In particular ...

Test #1: H_0: $P_{t,T} = Q_{t,T}$, the (observed) Risk-Neutral Distribution of Log-returns x

Test #2: H_0: $P_{t,T} = P_{HIST}$, a long-run Historical Distribution of Log-returns x

Equity-based estimation of CRRA γ

If $V_T = \widehat{S}_T$ (back to a total return index), the standard Q-martingale equity pricing relation is

$$\widehat{S}_t = \frac{1}{1+R^f_{t,T}} E^Q_t[\widehat{S}_T] \quad \text{or, equivalently,} \quad E^Q_t\left[R^e_{t,T} - R^f_{t,T}\right] = 0. \quad (*)$$

Fix T-t frequency (daily, monthly). Apply (*) on an unconditional basis [$p_t(x) \to p(x)$]:

$$\int e^{-\gamma x}\,(R^e_t - R^f_t)\,p(x_t)\,dx_t = 0, \quad \text{where} \quad R^e_t = e^{x_t} - 1 \quad [\, x_t = \text{log total return}\,]$$

For $p(x_t)$ use a long-run empirical density: $p(x) = \frac{1}{M}\sum_{i=1}^{M}\delta(x - x_i)$ [IID assumption]

Result: the <u>estimator</u> $f(\widehat{\gamma}) = 0$, where $f(\gamma) = \sum_{i=1}^{M} e^{-\gamma x_i}(e^{x_i} - 1 - R^f_i)$

This estimator is fast/easy (novel?) Just find a root. Bootstrap for Confidence Intervals.

Call it the "Equity IID γ-estimator". It uses only the empirical stock return series.

Equity-based estimation of γ (continued)

S&P 500 index monthly total returns; various start dates through Dec 2020

γ in [2,4] is plausible; $\gamma = 3.0 \pm 0.5$ is used for Covid-era ERP charts.

Similar results with CRSP data (Ken French data library; see next slide).

$\sigma(\gamma)$ from Monte Carlo bootstrap sampling.

My results: compatible with a classic study (Friend & Blume, 1975).

Finance literature has γ in [0,55]! (table in Bliss & Panigirtzoglou, 2004).

		CRRA		Ex Rets		Std Dev		Skewness		Kurtosis	
Start date	N_{obs}	$\hat{\gamma}$	$\sigma(\hat{\gamma})$	True	RN	True	RN	True	RN	True	RN
1. Jan 1926	1140	**2.32**	**0.59**	8.32	0	18.6	19.7	-0.50	-1.28	10.8	10.2
2. Jan 1950	852	**3.61**	**0.88**	7.97	0	14.5	15.4	-0.66	-0.94	5.4	6.2
3. Apr 1957	765	**3.03**	**0.89**	6.82	0	14.7	15.5	-0.68	-0.92	5.5	6.1
4. Jan 1987	408	**3.40**	**1.21**	8.54	0	15.3	16.8	-1.03	-1.39	6.5	7.0
5. Jan 1988	396	**3.87**	**1.27**	8.67	0	14.6	15.6	-0.75	-0.87	4.7	4.7

Equity-based estimation of γ (continued).

Fama-French CRSP data (all US-listed stocks; starts through Dec 2020).

Bootstrap error estimates $\sigma(\widehat{\gamma})$ (see next slide)

I. Daily Returns

Start date	N_{obs}	CRRA $\hat{\gamma}$	CRRA $\sigma(\hat{\gamma})$	Ex Rets True	Ex Rets RN	Std Dev True	Std Dev RN	Skewness True	Skewness RN	Kurtosis True	Kurtosis RN
1. Jul 1, 1926	24896	**2.58**	**0.59**	7.64	0	17.1	17.3	-0.48	-0.99	20.9	23.0
2. Jan 1, 1950	17953	**3.41**	**0.79**	7.72	0	15.2	15.5	-0.91	-1.62	23.8	30.9
3. Apr 1, 1957	16050	**2.83**	**0.81**	7.08	0	15.7	16.0	-0.87	-1.45	23.1	28.6
4. Jan 1, 1987	8570	**2.67**	**0.95**	9.12	0	18.4	18.8	-1.01	-1.59	22.0	26.6
5. Jan 1, 1988	8317	**2.90**	**0.96**	9.39	0	17.9	18.1	-0.51	-0.86	14.1	14.7

III. Monthly Returns

Start date	N_{obs}	CRRA $\hat{\gamma}$	CRRA $\sigma(\hat{\gamma})$	Ex Rets True	Ex Rets RN	Std Dev True	Std Dev RN	Skewness True	Skewness RN	Kurtosis True	Kurtosis RN
1. Jul 1, 1926	1134	**2.31**	**0.59**	8.18	0	18.5	19.6	-0.56	-1.21	9.7	9.4
2. Jan 1, 1950	852	**3.41**	**0.84**	8.02	0	14.9	16.0	-0.75	-1.03	5.7	6.4
3. Apr 1, 1957	765	**2.89**	**0.86**	7.03	0	15.3	16.2	-0.75	-0.99	5.6	6.2
4. Jan 1, 1987	408	**3.25**	**1.17**	8.70	0	15.8	17.4	-1.11	-1.38	6.8	7.3
5. Jan 1, 1988	396	**3.73**	**1.23**	8.93	0	15.1	16.2	-0.80	-0.91	4.8	4.7

Equity-based $\hat{\gamma}$ (continued)

The estimate here is based on Fama-French (F-F) daily CRSP returns from 1926-2024.

The data series is re-sampled 10,000 times. For each sample, the IID estimator is employed. The histogram shows the sampling distribution of the estimates.

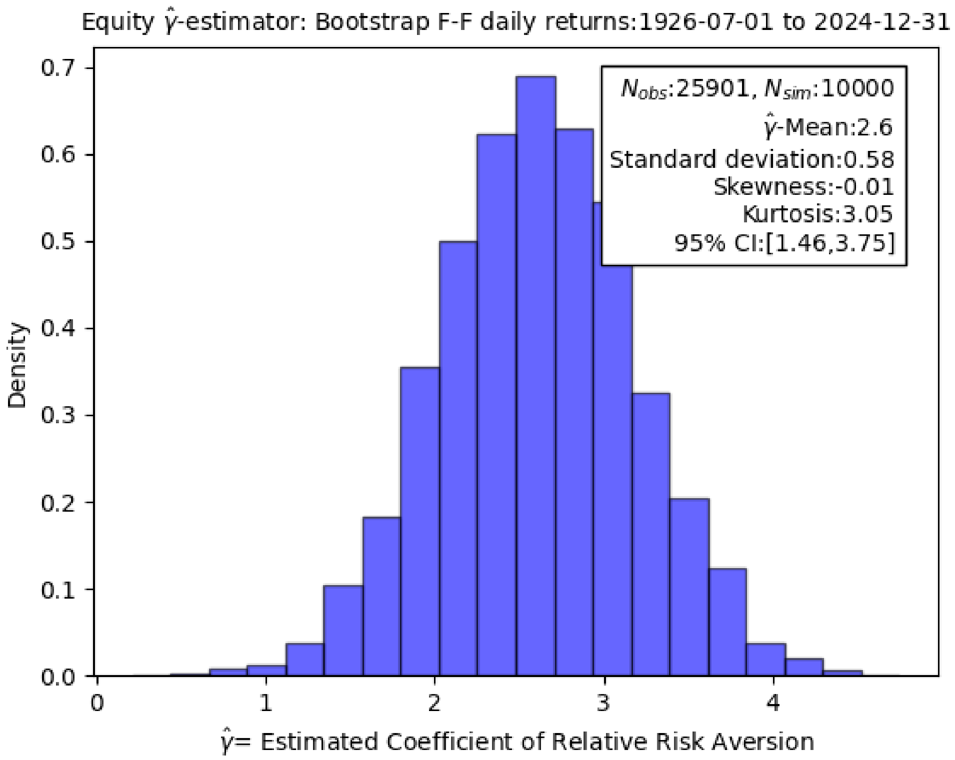

Equity-based $\hat{\gamma}$ (continued)

Is γ time-varying?

The Monte Carlo bootstrap is employed -- splitting (roughly 100 years) into four 25-year periods. The 95% confidence intervals for the γ estimates are shown. Also, the simple GBM equity estimator, based upon $\mu - r = \left(\gamma - \frac{1}{2}\right)\sigma^2$, is shown (crosses), which is virtually identical on daily returns. Here μ is the mean log total-return, r is the mean interest rate, and σ^2 is the variance of the log total-returns for each sample.

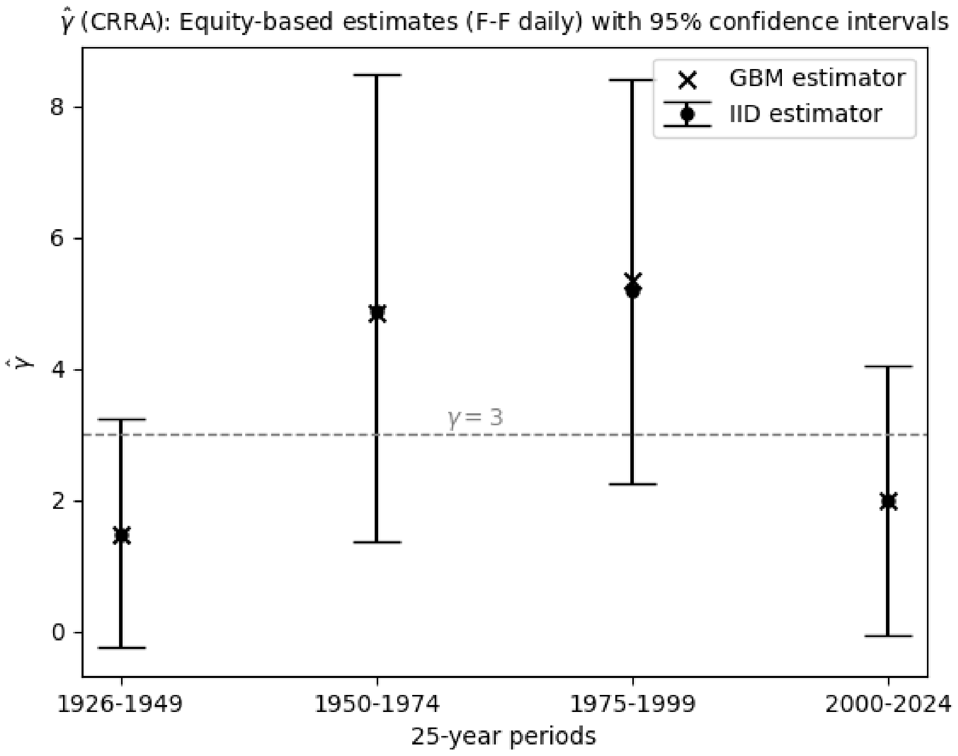

ERP term structures during the early days of the Covid-19 pandemic: $\gamma = 3 \pm 0.5$

Figure 10.1: **Timeline 1**

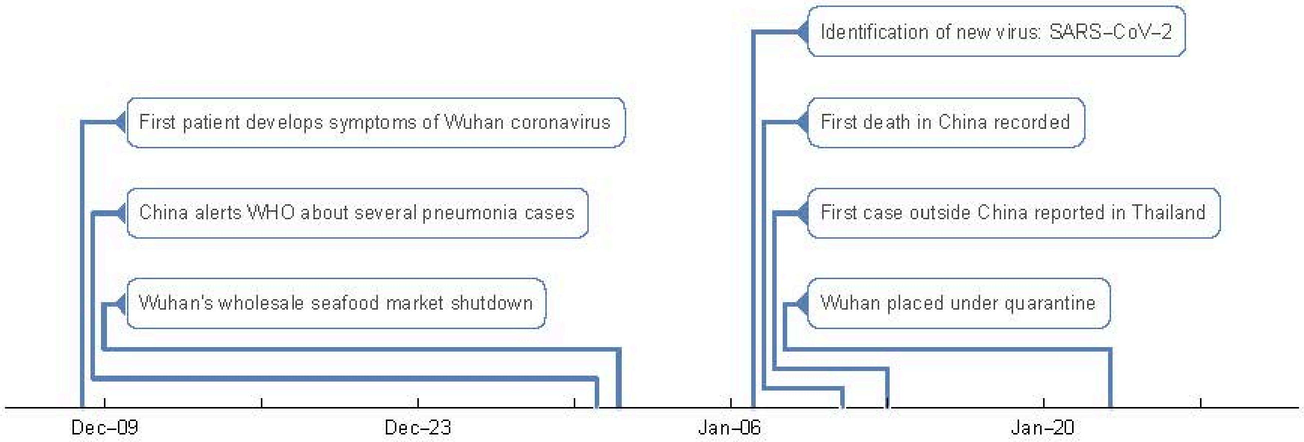

Figure 10.2: **US ERP term structure**

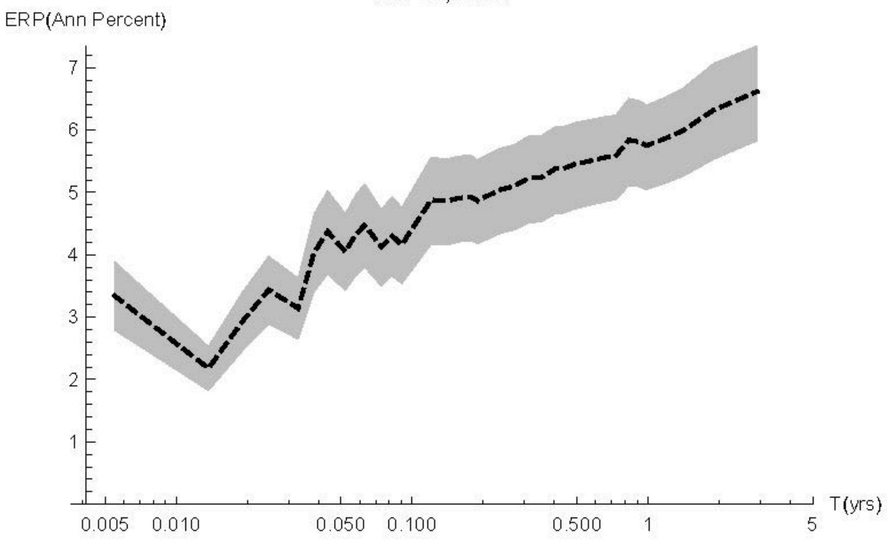

ERP term structures during the early days of the Covid-19 pandemic (continued)

Figure 10.7: **Timeline 4**

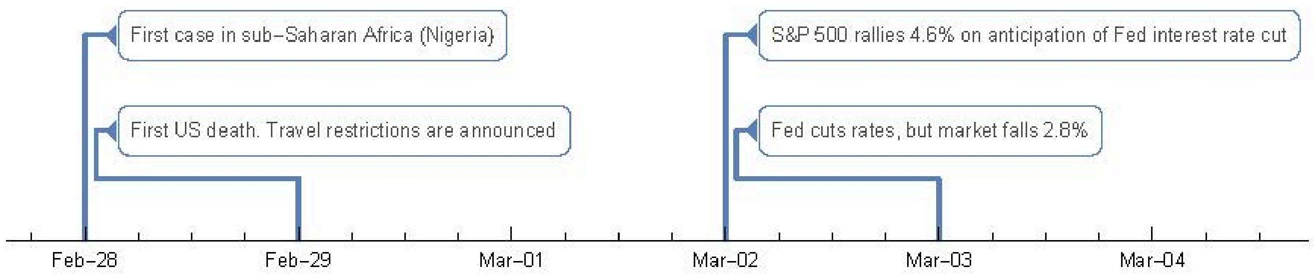

Figure 10.8: **US ERP term structure**

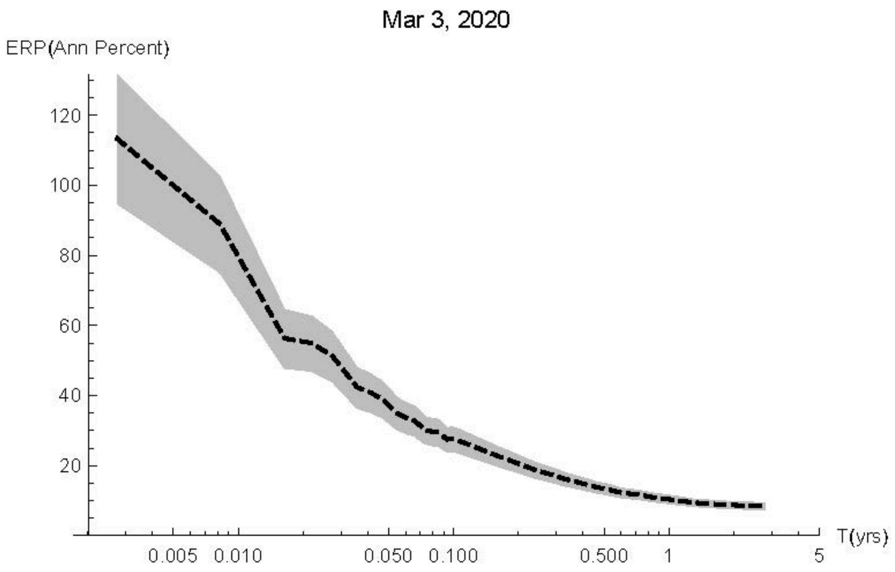

ERP term structures during the early days of the Covid-19 pandemic (continued)

Figure 10.9: **Timeline 5**

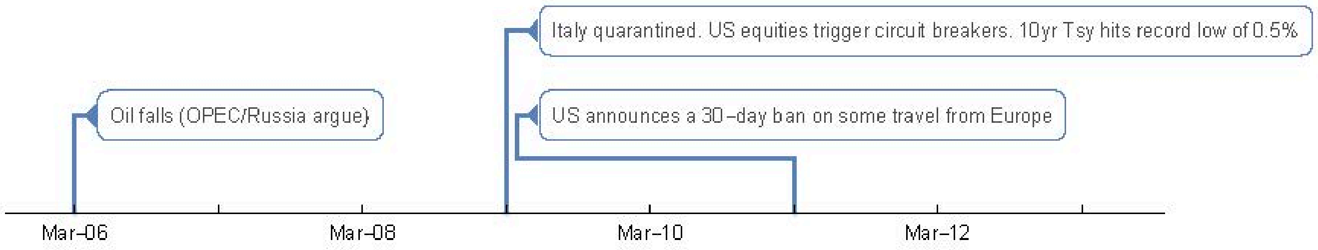

Figure 10.10: **US ERP term structure**

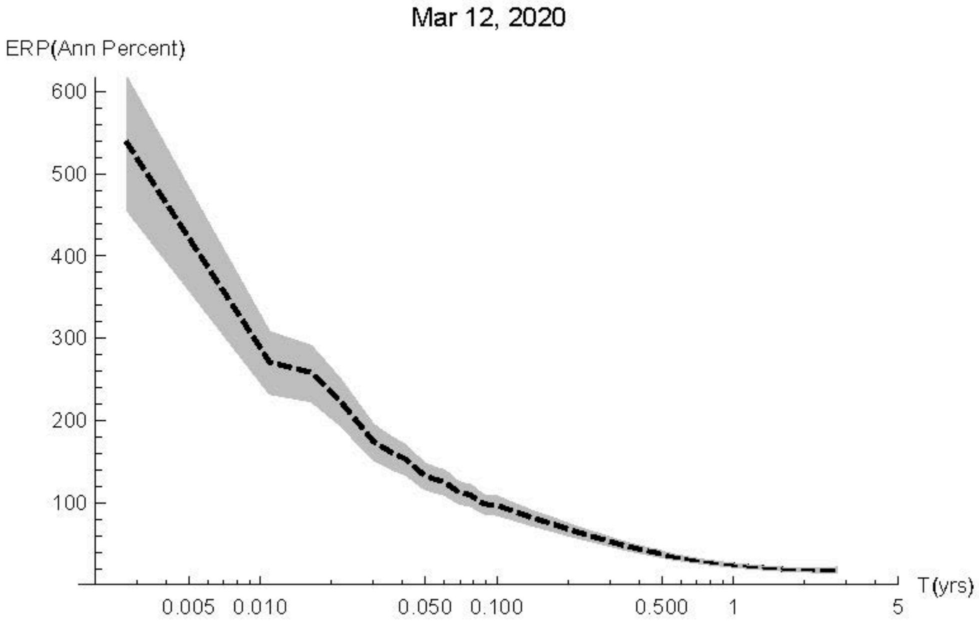

The Risk-return Trade-off

- Data: my Covid study, using GMM Q-fits.
- GMM = Gaussian mixture model.
- Dotted line shown is the (generalized) Merton's GBM relation: $\mathrm{ERP}_{t,T} = \gamma \times \left(\sigma^{P}_{t,T}\right)^2$,
- using γ = 3.

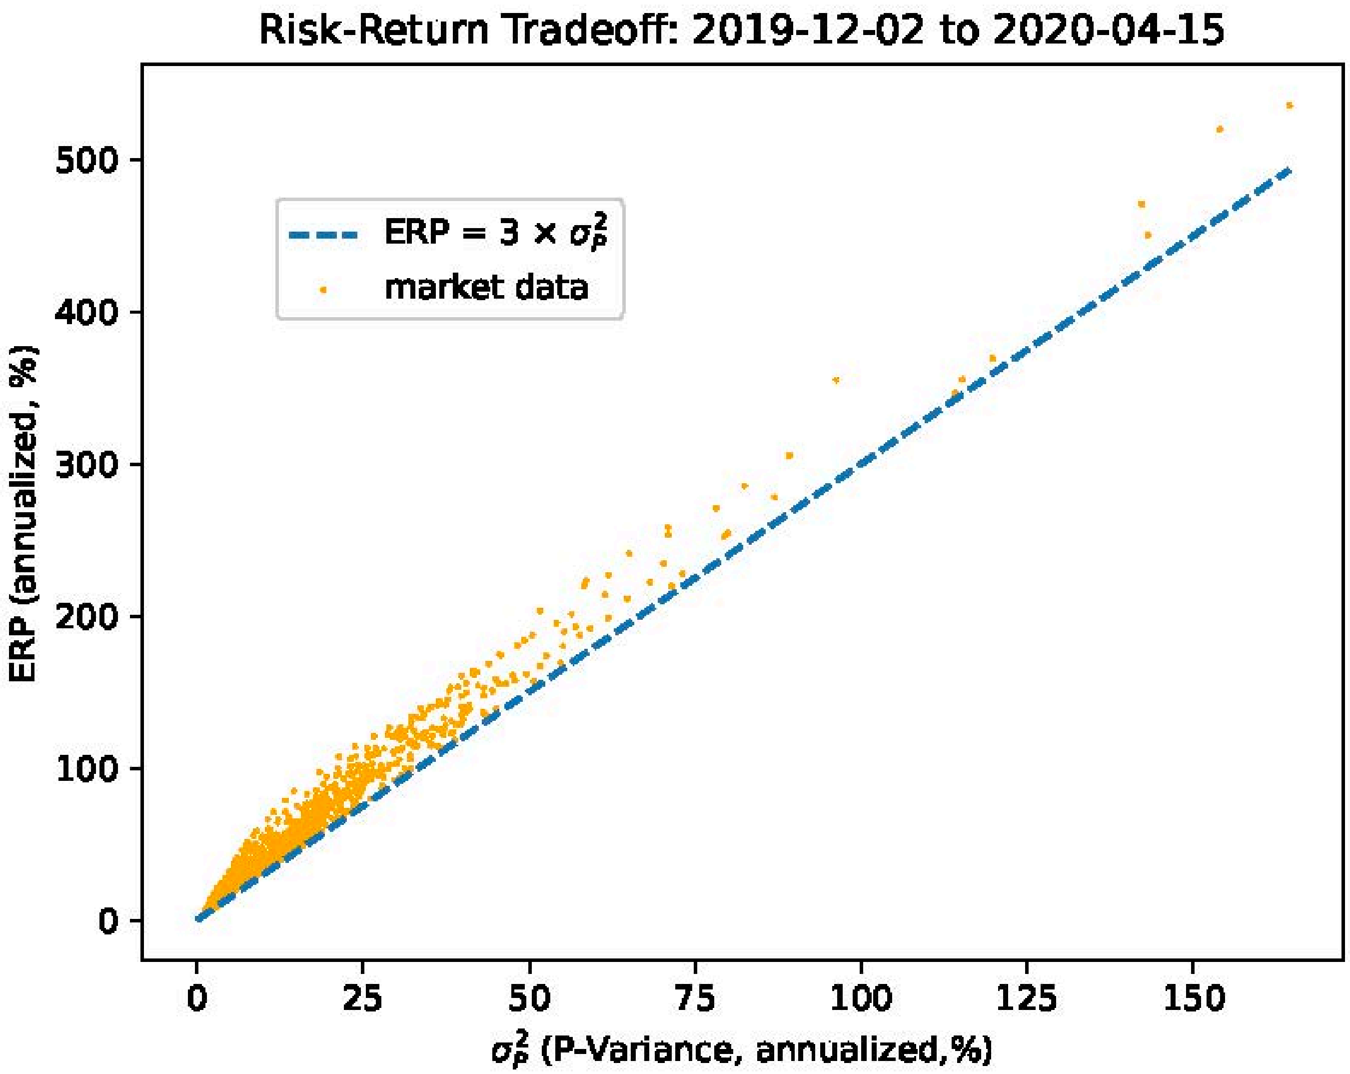

Estimating $\mathrm{ERP}_{t,T}$: Method 1: Fit a Gaussian Mixture Model (GMM) to Option Prices

- **This method: used for my Covid-era study and my Third-Friday study.**
- **With x_T = $\log S_T/S_t$ and $\tau = T - t$, N = 4 or 5,**
- $q_{X_T}(x) = \sum_{i=1}^{N} w_i \frac{e^{-(x-\mu_i\tau)^2/(2\,\sigma_i^2\tau)}}{\sqrt{2\,\pi\,\sigma_i^2\tau}}$ **(3N-2 free parameters; fit option prices)**
- **Abbreviate: $q \sim \mathrm{GMM}(w_i\ , \mu_i, \sigma_i^2)$**
- **After exponential tilt with γ, $p \sim \mathrm{GMM}(\widetilde{w}_i, \widetilde{\mu}_i, \sigma_i^2)$, where**
- $\widetilde{\mu}_i = \mu_i + \gamma\,\sigma_i^2$, $\widetilde{w}_i = \frac{w_i \exp(\beta_i)}{\sum_i w_i\ \exp(\beta_i)}$, **using** $\beta_i = \gamma\mu_i + \frac{1}{2}\gamma^2\ \sigma_i^2$

Estimating $\mathrm{ERP}_{t,T}$ after fitting a GMM (continued)

- **Using above GMM for $p_{X_T}(x)$; abbreviate div yield $\delta = \delta_{t,T}$, riskfree rate $r = r_{t,T}$,**
- $\Rightarrow \quad \mathrm{ERP}_{t,T}^{(ann,\%)} = \frac{100}{T-t}\left\{\left(e^{\delta\tau}\sum_{i=1}^{N}\tilde{w}_i e^{\alpha_i+(\gamma+\frac{1}{2})v_i}\right) - e^{r\tau}\right\}$
- **where** $\alpha_i = \mu_i\tau, \quad v_i = \sigma_i^2\tau, \quad \beta_i = \gamma\alpha_i + \frac{1}{2}\gamma^2 v_i,$
  $\tilde{w}_i = w_i e^{\beta_i} / \sum_{i=1}^{N} w_i e^{\beta_i}$

Fitting a GMM to option prices (continued)
The fitted prices often lie solidly between bid and asks

Figure 5: **Example of 'Good' GMM fit:**
Trade date: Feb 7, 2018. Expiration date: Feb 12, 2018

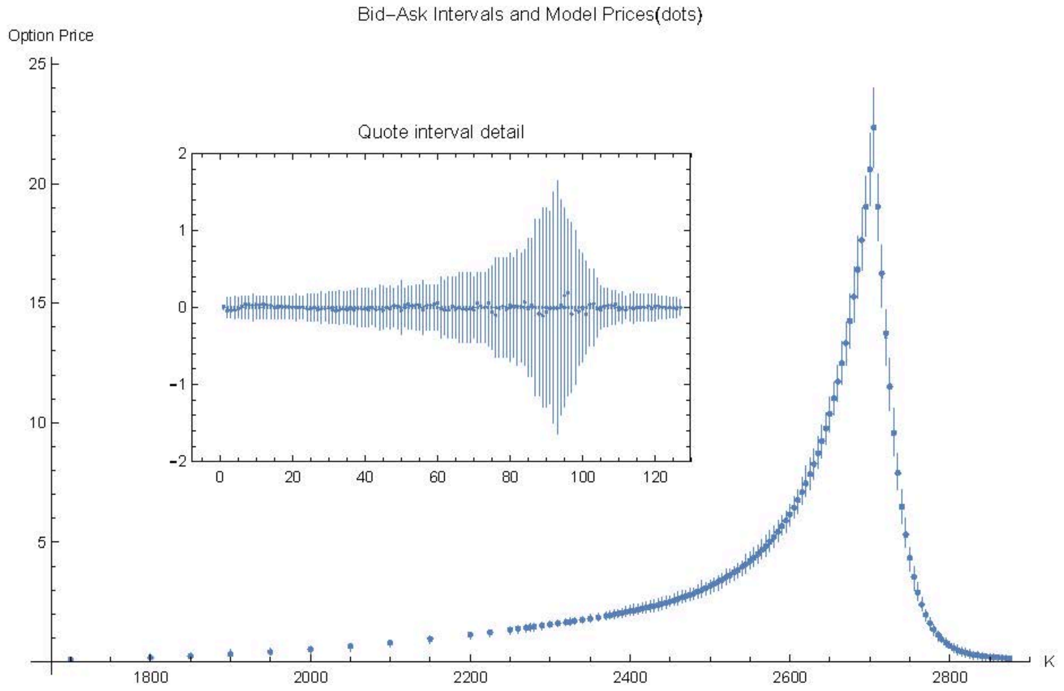

Two option chain methods:
1. GMM (continued)

Typical results from fitting the Risk-neutral (Q-density) to SPX options, and then tilting by γ to find the Real-world (P-density).

Columns: Q,P, and Q-P densities (almost always find unimodal fits)

Figure 12: **RNDs (Q) and RWDs (P): Feb 7, 2018 (cont.)**

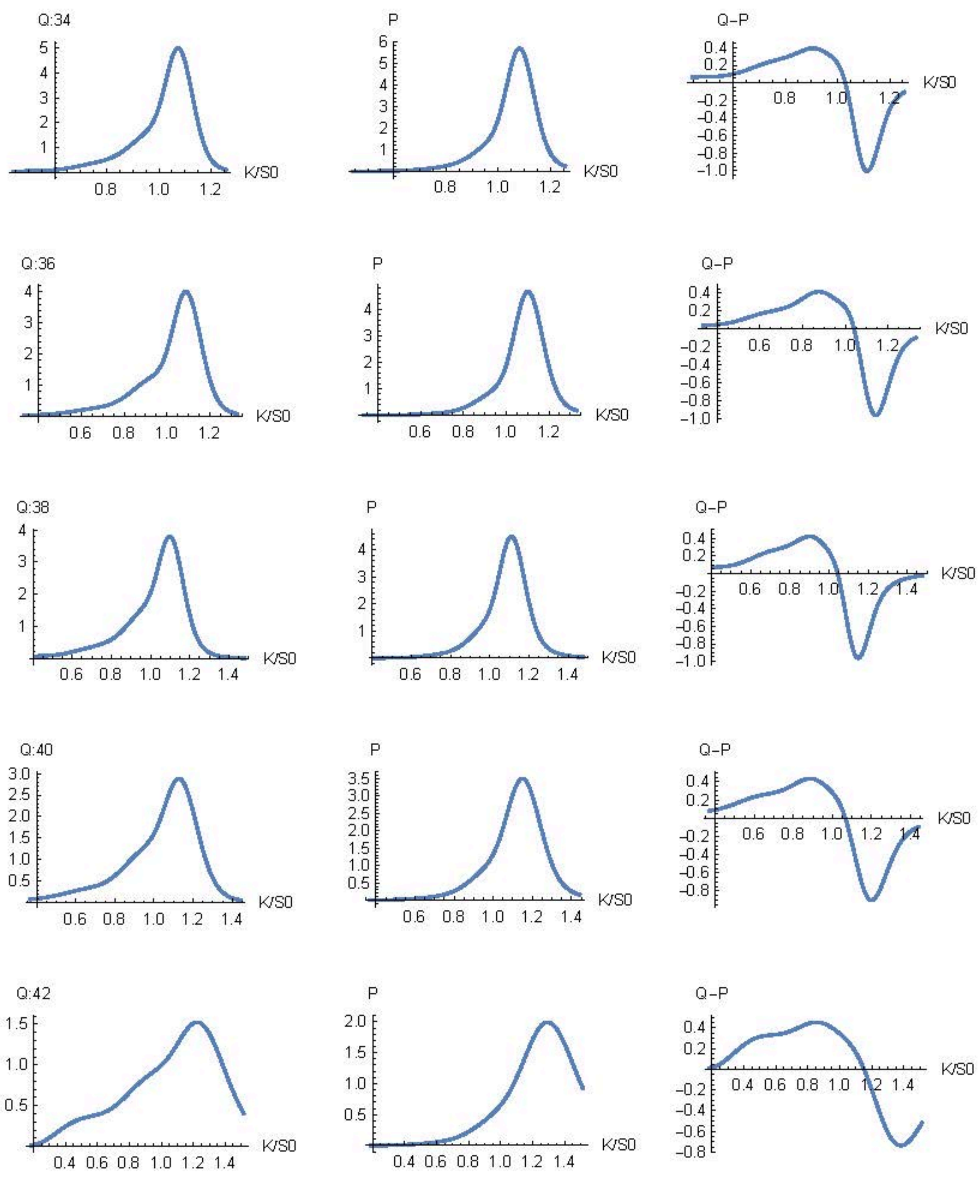

Fitting a GMM to option prices (continued).

Fitted GMM smiles are smooth, arbitrage-free, close to the market and asymptotic to a <u>constant</u> implied volatility.

Figure 7: **Smile and Extended Smile Fits:**
Trade date: Feb 7, 2018. Expiration date: Feb 12, 2018

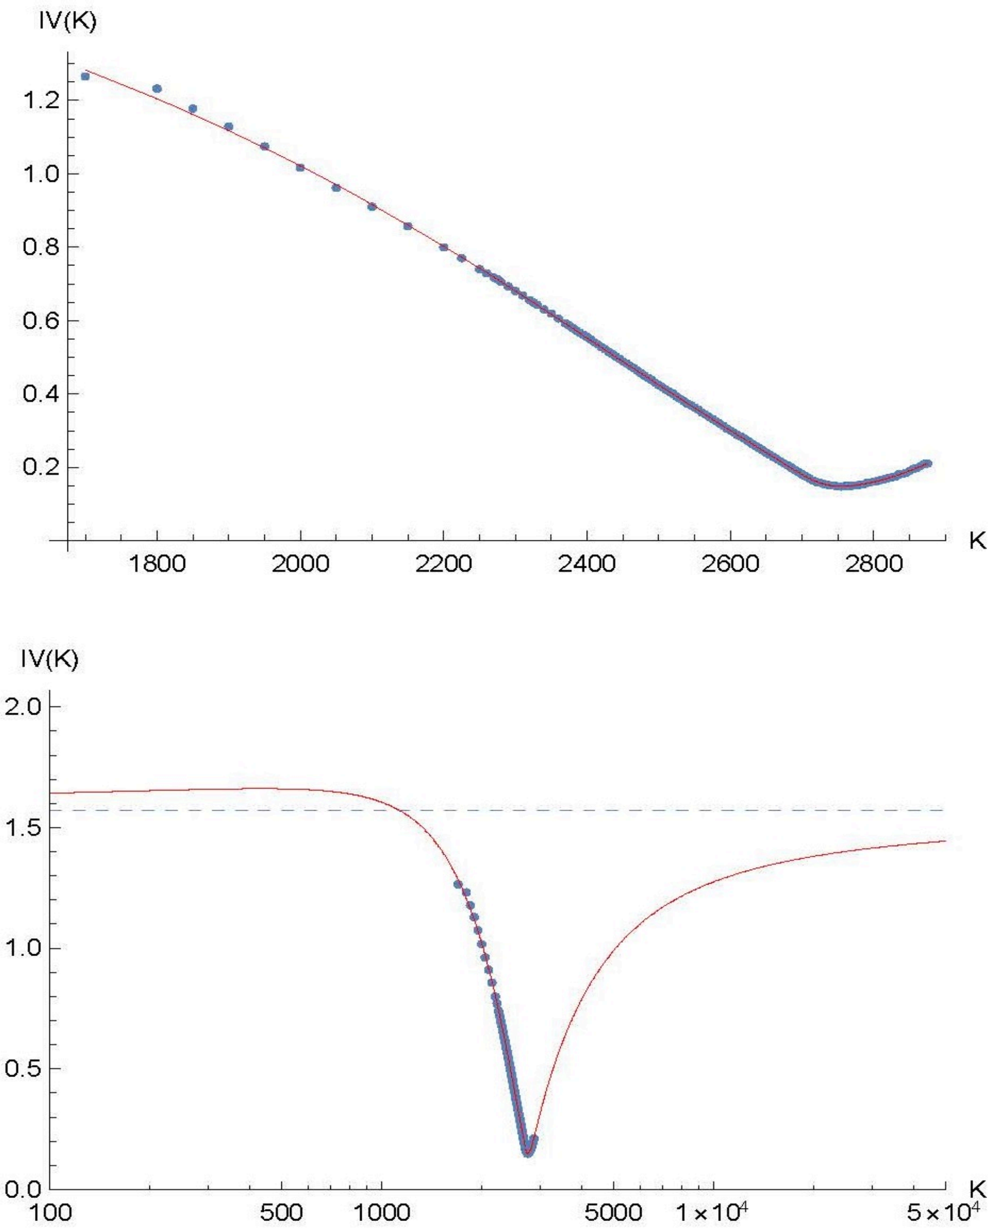

Estimating $\mathrm{ERP}_{t,T}$: option-based method #2. VIX Technology (Payoff replication)

- Recall that the CBOE Volatility Index ("VIX") uses SPX options to estimate:
- $VIX_{0,T}^2 = E_0^Q\left[\frac{1}{T}\int_0^T \sigma_t^2\,dt\right] = -\frac{2}{T}E_0^Q\left[\log\frac{S_T}{F_{0,T}}\right]$, where σ_t^2 is an inst. variance rate, $F_{0,T}$ = forward stock price for sale at T, S_T = stock price (no dividends).
- Assume a continuum of puts and calls with prices $P_{0,T}(K), C_{0,T}(K)$, (K=strike).
- From Breeden - Litzenberger ($Q_{0,T}(K) = e^{rT}C''_{0,T}(K)$) and parts, value a general payoff $w(S_T)$ using the (well-known) general replication formula:
- $E_0^Q[w(S_T)] = w\left(F_{0,T}\right) + e^{rT}\left\{\int_0^{F_{0,T}} w''(K)\,P_{0,T}(K)dK + \int_{F_{0,T}}^{\infty} w''(K)\,C_{0,T}(K)dK\right\}$.
- In practice, use options-implied forward (See "VIX White Paper").
- Notational abuse here: $P_{0,T}(K)$ is the put option price at strike K.

VIX Technology (continued)

- **Recall** $\mathrm{ERP}_{0,T}$ **=** $\frac{100}{T}\left\{E_0^P\left[\frac{\widehat{S}_T}{\widehat{S}_0}\right]-\left(1+\ R_{0,T}^f\right)\right\}$**, where** $\widehat{S}_t$ **= price (<u>with</u> dividends).**
- **<u>Want</u> Q-expectations, using density** $q_{S_T}(s)$, **where** S_t **= price (<u>no</u> dividends).**
- **With** $\widehat{S}_T$ **=** $e^{\delta T}S_T$**, where** $\delta=\delta_{0,T}$ **is the (options-implied) dividend yield, apply the exponential tilt** $\mathrm{P}\rightarrow Q$**. The result allows application of the general replication formula (previous slide):**
- $$E_0^P\left[\frac{\widehat{S}_T}{\widehat{S}_0}\right]=\frac{e^{\delta T}}{S_0}\frac{E_0^Q[w_{num}(S_T)]}{E_0^Q[w_{den}(S_T)]}=\frac{e^{\delta T}}{S_0}\frac{E_0^Q\left[S_T^{1+\gamma}\right]}{E_0^Q\left[S_T^{\gamma}\right]}$$

ERP via VIX Tech (continued)

Using the replication formula, one finds the boxed result (K_0 is an arbitrary positive strike).

Then, integrals $\rightarrow$ discrete sums as in VIX white paper.

Note: NO OPTIMIZATION, so extremely fast!

$$\boxed{\begin{aligned}
&\mathrm{ERP}_{0,T} = \mathrm{e}^{rT}\left\{\left(\frac{K_0}{F_{0,T}}\right)\left(\frac{N}{D}\right) - 1\right\}, \quad \text{using} \\
&N = 1 + (1+\gamma)\left(\frac{F_{0,T}}{K_0} - 1\right) + \frac{\gamma(1+\gamma)}{K_0^{1+\gamma}}\,\mathrm{e}^{rT}\, I_1, \\
&D = 1 + \gamma\left(\frac{F_{0,T}}{K_0} - 1\right) + \frac{\gamma(\gamma-1)}{K_0^{\gamma}}\,\mathrm{e}^{rT}\, I_2, \\
&I_1 = \left\{\int_0^{K_0} P_{0,T}(K)\,K^{\gamma-1}\,dK + \int_{K_0}^{\infty} C_{0,T}(K)\,K^{\gamma-1}\,dK\right\}, \\
&I_2 = \left\{\int_0^{K_0} P_{0,T}(K)\,K^{\gamma-2}\,dK + \int_{K_0}^{\infty} C_{0,T}(K)\,K^{\gamma-2}\,dK\right\},
\end{aligned}}$$

Note δ and S_0 dependencies have disappeared in favor of just r and $F_{0,T}$.

Think of the denominator expression above as a function of γ: $D =$ then the numerator $N(\gamma) = D(1+\gamma)$.

As a simple check, suppose a risk-neutral world: the representative has CRRA $\gamma = 0$. Then (13.7) yields

$$\mathrm{ERP}_{0,T} = \mathrm{e}^{rT}\{1 - 1\} = 0, \quad \text{as expected.}$$

Table 13.1: **VWP-ERP vs. GMM-ERP method (Feb 7, 2018)**
Results for ERP's, comparing the (fast) VWP-ERP method vs. the earlier GMM-ERP method of (Lewis, 2019). On dual expirations, only the SPXW entry is taken from (Lewis, 2019). The VWP method is essentially instantaneous relative to the GMM method, generating this table in less than a second of computer time. The last column shows the VIX as reported by the Cboe at approximately 15:45 EST, 15 minutes from the regular close. Boldface provides a guide to the eye for ERP compares.

File Num	Option Exp Date	days	T	VWP method r	VWP method q	VWP method forward	VWP method myVIX	VWP-ERP	GMM-ERP	Cboe VIX
1	02-09-2018	2	0.005	1.24	10.88	2705.1	33.02	**26.3**	**26.3**	
2	02-12-2018	5	0.014	1.25	6.05	2704.7	26.66	**16.0**	**16.2**	
3	02-14-2018	7	0.019	1.26	5.85	2704.1	26.77	**16.0**	**16.2**	
4	02-16-2018	9	0.025	1.27	5.66	2703.5	27.33	**15.9**	**15.9**	
5	02-20-2018	13	0.036	1.28	3.60	2704.2	23.44	**12.1**	**12.3**	
6	02-21-2018	14	0.038	1.29	3.20	2704.5	23.61	**12.3**	**12.4**	
7	02-23-2018	16	0.044	1.30	2.76	2704.7	24.89	**12.7**	**12.8**	
8	02-26-2018	19	0.052	1.31	2.79	2704.4	22.76	**11.3**	**11.3**	
9	02-28-2018	21	0.058	1.32	2.72	2704.3	24.09	**11.6**	**11.8**	
10	03-02-2018	23	0.063	1.32	2.87	2703.8	24.03	**11.6**	**12.1**	
11	03-05-2018	26	0.071	1.34	2.52	2704.2	21.97	**10.6**	**10.9**	
12	03-07-2018	28	0.077	1.34	2.37	2704.3	21.99	**10.6**	**10.9**	
13	03-09-2018	30	0.082	1.35	2.35	2704.2	23.37	**10.7**	**11.5**	
14	03-12-2018	33	0.090	1.36	2.21	2704.4	21.48	**9.9**	**10.2**	
15	03-14-2018	35	0.096	1.37	2.36	2703.9	21.53	**9.9**	**10.2**	
16	03-16-2018	37	0.101	1.38	2.32	2703.9	22.69	**10.1**	**10.4**	
17	03-23-2018	44	0.121	1.40	1.78	2705.3	21.96	**9.6**	**10.0**	
18	03-29-2018	50	0.137	1.42	1.78	2705.2	21.91	**9.3**	**9.6**	
19	04-06-2018	58	0.159	1.45	1.55	2706.1	21.21	**8.8**	**8.9**	
20	04-20-2018	72	0.197	1.49	1.47	2706.6	21.44	**8.6**	**8.6**	
21	04-30-2018	82	0.225	1.52	1.33	2707.6	21.06	**8.3**	**8.4**	
22	05-18-2018	100	0.274	1.56	1.51	2706.9	21.08	**8.3**	**8.5**	
23	05-31-2018	113	0.310	1.59	1.55	2706.8	20.93	**8.1**	**8.3**	
24	06-15-2018	128	0.351	1.62	1.55	2707.2	20.94	**8.0**	**8.2**	
25	06-29-2018	142	0.389	1.65	1.45	2708.6	20.65	**7.9**	**8.1**	
26	07-31-2018	174	0.477	1.70	1.34	2711.2	20.09	**7.6**	**7.8**	
27	09-21-2018	226	0.619	1.77	1.39	2712.9	20.16	**7.5**	**7.5**	
28	09-28-2018	233	0.638	1.78	1.40	2713.1	19.99	**7.4**	**7.5**	
29	12-21-2018	317	0.868	1.86	1.34	2718.8	19.94	**7.3**	**7.4**	
30	12-31-2018	327	0.896	1.86	1.26	2721.3	19.60	**7.1**	**7.5**	
31	01-18-2019	345	0.945	1.88	1.27	2722.0	19.55	**7.1**	**7.3**	
32	03-15-2019	401	1.099	1.92	1.33	2724.0	19.76	**7.1**	**7.3**	
33	06-21-2019	499	1.367	1.98	1.36	2729.6	19.88	**7.1**	**7.3**	
34	12-20-2019	681	1.866	2.10	1.40	2741.8	19.76	**7.0**	**7.2**	
35	12-18-2020	1045	2.863	2.28	1.44	2772.8	20.26	**7.2**	**7.7**	

My two ERP methods are numerically very close

Do the ERP's Predict the Market Return? Only weakly

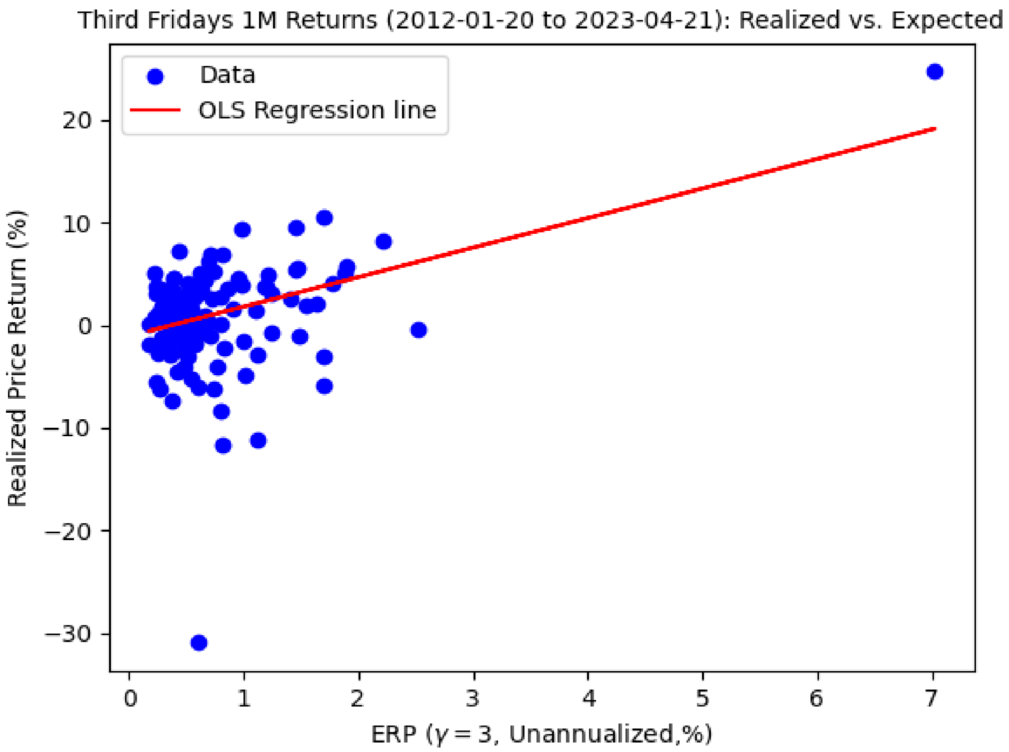

Do the ERP's Predict the Market Return? Only weakly

```
                            OLS Regression Results
==============================================================================
Dep. Variable:                      y   R-squared:                       0.163
Model:                            OLS   Adj. R-squared:                  0.157
Method:                 Least Squares   F-statistic:                     4.884
Date:                Tue, 18 Feb 2025   Prob (F-statistic):             0.0288
Time:                        06:08:02   Log-Likelihood:                -401.03
No. Observations:                 136   AIC:                             806.1
Df Residuals:                     134   BIC:                             811.9
Df Model:                           1
Covariance Type:                  HC3
==============================================================================
                 coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
const         -1.0810      0.917     -1.179      0.239      -2.879       0.717
x1             2.8783      1.302      2.210      0.027       0.326       5.431
==============================================================================
Omnibus:                      113.121   Durbin-Watson:                   2.083
Prob(Omnibus):                  0.000   Jarque-Bera (JB):             1445.693
Skew:                          -2.806   Prob(JB):                         0.00
Kurtosis:                      17.954   Cond. No.                         2.44
==============================================================================

Notes:
[1] Standard Errors are heteroscedasticity robust (HC3)
```

The Key Issue: What is the best systematic forecast of $P_{t,T}(x)$?

My proposal is $P_{t,T}^{(\gamma)}(x) = C\, e^{\gamma x}\, Q_{t,T}(x)$ for some $\gamma > 0$.

What are some plausible alternatives; i.e., Null Hypotheses H_0?

H_0 #1: Perhaps $P_{t,T}(x) = Q_{t,T}(x)$ (the option-implied risk-neutral distribution)

H_0 #2: Perhaps $P_{t,T}(x) = P_{hist}(x)$ (a fitted long-run historical distribution)

- We can test Alternative #1 by a Maximum Likelihood Analysis. From my Third Friday study, I have GMM (N=4) fits to $Q_{t,T}(x)$ for each of 136 trade dates, where t is a monthly third Friday and T is the subsequent third Friday.
- I have the realized log-price-returns $\boldsymbol{x} = \{x_1, x_2, \ldots, x_{136}\}$.
- Suppose the market return Data Generating Process each month was an independent draws from $P_{t,T}^{(\gamma)}(x)$.

The Key Issue: What is the best systematic forecast of $P_{t,T}(x)$? (continued)

Suppose the market return Data Generating Process each month was an independent draw from $P_{t,T}^{(\gamma)}(x)$ for some γ (fixed for all the months).

Then, the Likelihood of seeing the market data is:

$$\text{Likelihood}(x) = P_1^{(\gamma)}(x_1)\ P_2^{(\gamma)}(x_2) \cdots\ P_{136}^{(\gamma)}(x_{136}).$$

First, we find a new options-based γ-estimator:

$$\gamma^* = \text{argmax}_\gamma \{\text{Log Likelihood }(\text{x})\} = \text{argmax}_\gamma \sum_{i=1}^{136} \log P_i^{(\gamma)}(x_i)$$

Next slide shows the result for γ^*...

The Key Issue: What is the best systematic forecast of $P_{t,T}(x)$? (cont)

Optimal γ from Maximum Likelihood of Observed Market Returns

Confidence interval from Wilks' Theorem here, which involves critical values of a χ^2-distribution(df = 1).

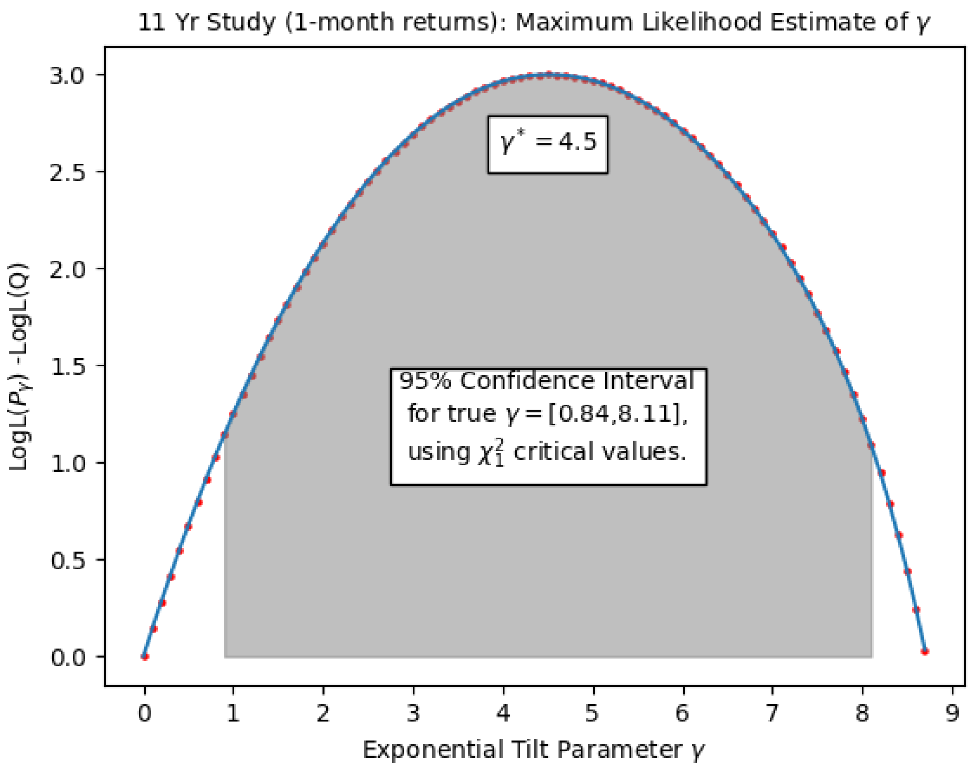

Recall Wilks' Theorem (1938)

Simplest version. We have a one-parameter data-generating probability density P_θ(x), where the unknown parameter θ is to be estimated by Maximum Likelihood. In our case, $\theta = \gamma$, the CRRA.

To do so, given a sequence of n independent observations $\{x_1, x_2, \ldots, x_n\}$, we form the estimate:

$$\theta^* = \text{argmax}_\theta \{\text{Log Likelihood } (\mathbf{x})\} = \text{argmax}_\theta \sum_{i=1}^{n} \log P_\theta(x_i).$$

A competing alternative is the Null Hypothesis H_0 that the true $\theta = 0$. Wilks proved that the Likelihood Ratio Test Statistic (LRT) has the following asymptotic behavior:

$$\text{LRT} \equiv 2\left\{\sum_{i=1}^{n} \log P_{\theta^*}(x_i) - \sum_{i=1}^{n} \log P_0(x_i)\right\} \to \chi^2 \text{ – distribution (1 degree freedom),}$$

as $n \to \infty$, **when H_0 is true.**

Note: f_{χ^2}(x) = $\frac{e^{-x/2}}{\sqrt{2\pi x}}$ is simply the pdf of the square of a N(0,1) variate.

The Key Issue: What is the best systematic forecast of $P_{t,T}(x)$? (continued)

Monte Carlo Bootstrap Distribution of Likelihood Ratio Test Statistic (LRT).

If the Null H_0 is true, P=Q: simulated returns should be drawn from GMM Q-distributions.

Shows the (Wilks' Theorem) χ^2-distribution is a reasonable LRT limit.

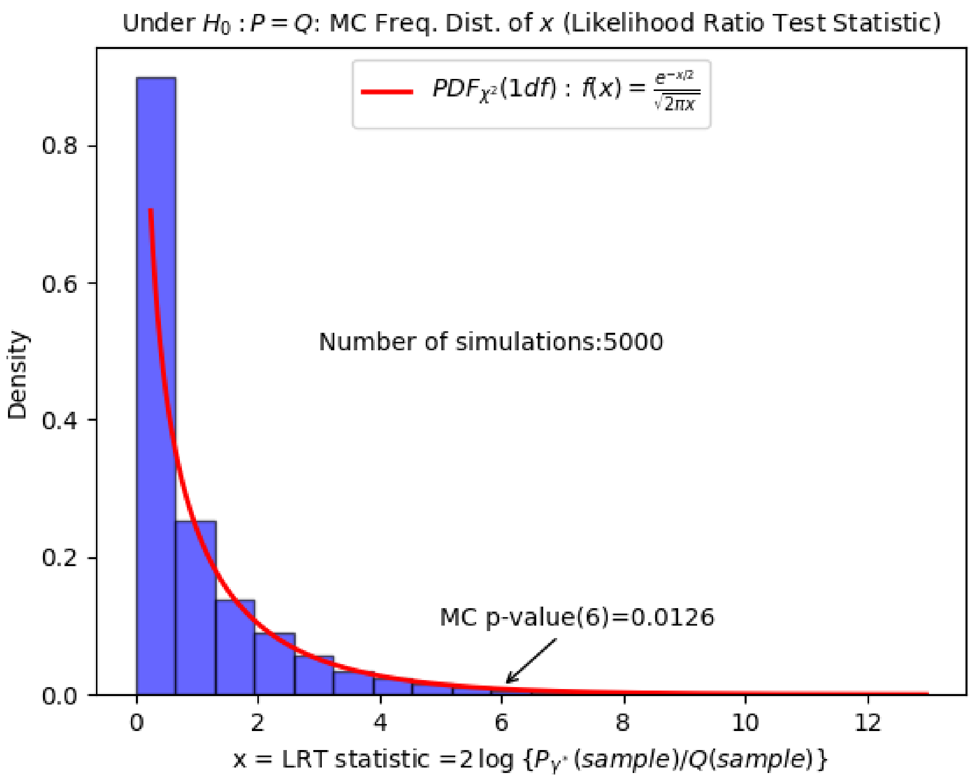

Data and automation notes

- My codes started as a mix of Mathematica (GMM) and Python (VIX tech). At this point they are entirely Python, with a lot of help from Microsoft Copilot. Python GMM fits use scipy.optimize.minimize (method = 'SLSQP')
- Option data from the Cboe DataShop (modest pricing).
- Cboe offers: 1. End of day option quotes with calcs (3:45pm EST snapshot)
- 2. EOD quotes: 3:45pm EST + 4:15pm EST (SPX option close).
- I use/recommend the 3:45pm quotes (high liquidity; 4pm=NYSE close).
- The 3:45pm option data is available from 2012 to date.

Conclusions

- The ERP has substantial time and term structure variation.
- That variation is best seen through option data and Q→P transformations.
- The exponential tilt transformation is a straightforward, plausible, and economically motivated, one-parameter, dynamics-free model.
- Competing Hypotheses, that the true P-data generating probability distributions are either: (i) the option-implied Q-distribution, or (ii) a long-run historical P-distribution, are rejected at 95% Confidence levels.
- ERP term structures are estimated quickly with VIX-type methods.
- Option-implied Q-distributions are readily estimated by fitting Gaussian Mixtures.
- Equity-based and options-based γ-estimators yield similar results.

A few selected references

- I. Friend and M.E. Blume, The Demand for Risky Assets, Amer. Econ. Review, 65, No. 5, (Dec 1975).
- F. Duarte and C. Rosa, The Equity Risk Premium: A Review of Models, (staff report) NY Fed (2015).
- R.R. Bliss and N. Panigirtzoglou, Option-implied risk aversion estimates, J. Finance, 59(1), (2004).
- R.C. Merton, On Estimating the Expected Return on the Market, J. Financial Economics, 8(4), 323-361, (1980) [See Appendix A]
- A.L. Lewis, The Equity Risk Premium: An Options-based Approach, forthcoming book, Finance Press, (in the meantime, see arXiv:1910.14522 and arXiv:2004.13871).