Simple Variance Swaps

Ian Martin
ian.martin@stanford.edu

LSE/Stanford and NBER

December 7, 2012

Outline

- Variance swaps and simple variance swaps
- VIX and SVIX
- A lower bound on the equity premium

A lower bound on the equity premium

1 year horizon, in %

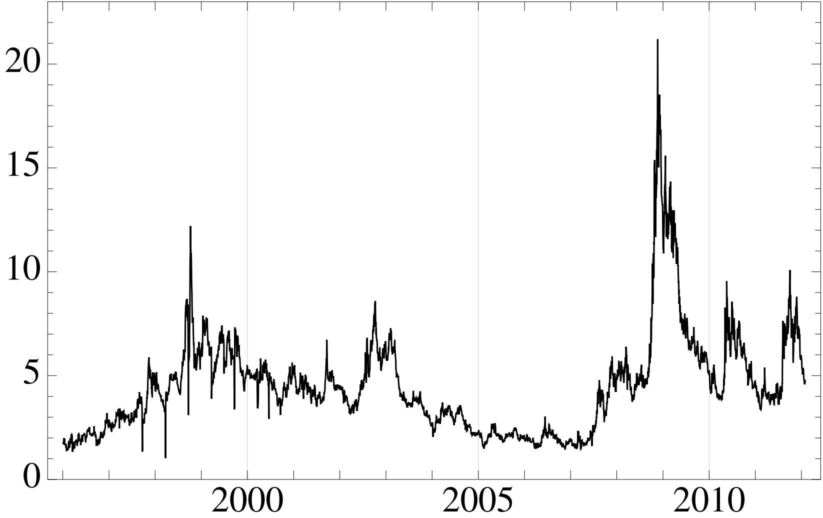

The equity premium

According to John Campbell's Princeton Lecture in Finance

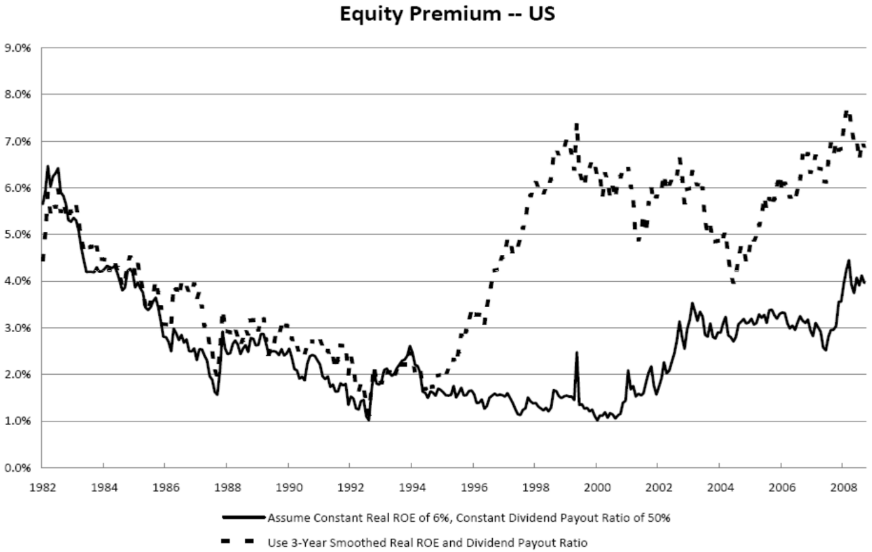

A lower bound on the equity premium

1 month horizon, annualized, in %

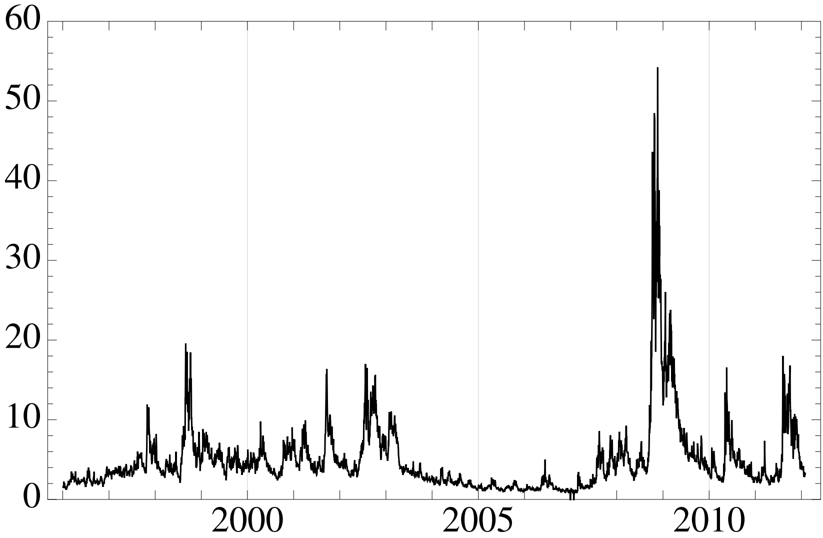

Assumptions

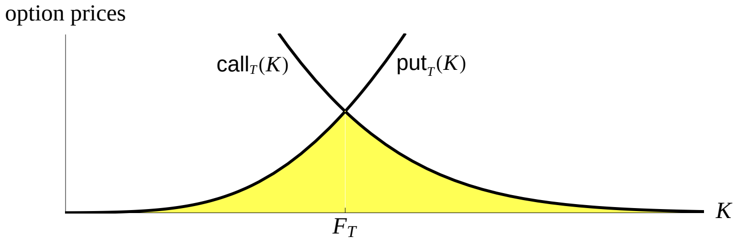

- No arbitrage
- Perfectly liquid market in options
- Underlying asset does not pay dividends (partially relaxed in paper)
- Constant riskless rate between now and time T (next few slides only)

Alternative ways to measure variance

- In a Black–Scholes world, straddles measure variance of log returns
  - How VIX used to be measured
- Once we acknowledge that volatility is time-varying, it is natural to look at variance swaps, which measure expected time-averaged instantaneous variance of log returns (. . . or do they?)
  - How VIX is now measured
- Once we also acknowledge that there are jumps, it is natural to look at simple variance swaps, which measure variance of simple returns
  - How VIX should be measured?

Ways to measure variance (1): straddles

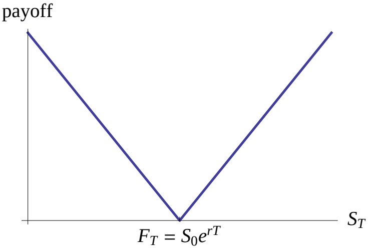

- In Black–Scholes world, any option price reveals volatility, σ
- *At-the-money-forward straddles* are particularly convenient because the price of an ATMF straddle $\approx 0.8 \cdot S_0\sigma\sqrt{T}$
  - To be precise, straddle price $= \sqrt{2/\pi} \cdot S_0\sigma\sqrt{T} \cdot (1 + \text{error})$, where $|\text{error}| < \sigma^2 T/24$
  - If $\sigma = 20\%$ and $T = 1/12$ then $|\text{error}| < 0.00014$

Ways to measure variance (2): variance swaps

- A *variance swap* is an agreement to exchange

$$\left(\log\frac{S_\Delta}{S_0}\right)^2+\left(\log\frac{S_{2\Delta}}{S_\Delta}\right)^2+\cdots+\left(\log\frac{S_T}{S_{T-\Delta}}\right)^2$$

  for a pre-agreed "strike" $\widetilde{V}$, at time T
- $\widetilde{V}$ is chosen so that no money needs to change hands up front:

$$\mathbb{E}^*\left[\left(\log\frac{S_\Delta}{S_0}\right)^2+\left(\log\frac{S_{2\Delta}}{S_\Delta}\right)^2+\cdots+\left(\log\frac{S_T}{S_{T-\Delta}}\right)^2-\widetilde{V}\right]=0$$

- If the underlying asset price follows *any* diffusion, then as $\Delta\to 0$, we end up with $\widetilde{V}=\mathbb{E}^*\left[\int_0^T\sigma_t^2\,dt\right]$
- Asterisks indicate risk-neutral expectations etc. $\mathbb{E}^*(X)=e^{rT}\,\mathbb{E}(MX)$

Ways to measure variance (2): variance swaps

- Exploiting Itô's lemma,

$$
\begin{aligned}
\widetilde{V} &= \mathbb{E}^* \left[\int_0^T \sigma_t^2 \, dt \right] \\
&= 2\,\mathbb{E}^* \left[\int_0^T \frac{1}{S_t} \, dS_t - \int_0^T d \log S_t \right] \\
&= 2rT - 2\,\mathbb{E}^* \log \frac{S_T}{S_0}
\end{aligned}
$$

- We could price the variance swap if a "log security" were traded (Neuberger 1994)
- Using Breeden–Litzenberger 1978 logic, can work out price of log security from European option prices (Carr–Madan 1998)

Ways to measure variance (2): variance swaps

- $\widetilde{V}$ is calculated from option prices

$$\widetilde{V} = 2e^{rT} \left\{ \int_0^{F_T} \frac{1}{K^2} \, \text{put}_T(K) \, dK + \int_{F_T}^{\infty} \frac{1}{K^2} \, \text{call}_T(K) \, dK \right\}$$

- Hedge:
  - (i) a static position in options expiring at T: hold OTM puts and OTM calls in quantities proportional to $1/K^2$
  - (ii) a dynamic delta-hedge
- VIX is calculated from this formula
- The pricing and hedging of variance swaps is often referred to as model-free

Variance swap pricing is not model-free

The cataclysm that hit almost all financial markets in 2008 had particularly pronounced effects on volatility derivatives . . . . Dealers learned the hard way that the standard theory for pricing and hedging variance swaps is not nearly as model-free as previously supposed . . . . In particular, sharp moves in the underlying highlighted exposures to cubed and higher-order daily returns. . . . This issue was particularly acute for single names, as the options are not as liquid and the most extreme moves are bigger. As a result, the market for single-name variance swaps has evaporated in 2009.

—Carr & Lee, *Annual Review of Financial Economics*, 2009

Ways to measure variance (3): simple variance swaps

- A *simple variance swap* is an agreement to exchange

$$\left(\frac{S_\Delta - S_0}{F_0}\right)^2 + \left(\frac{S_{2\Delta} - S_\Delta}{F_\Delta}\right)^2 + \cdots + \left(\frac{S_T - S_{T-\Delta}}{F_{T-\Delta}}\right)^2$$

  for a pre-agreed strike V, at time T

- F_t is the forward price of the underlying to time t, known at time 0
- Any constants known at time 0 could be put in the denominator, but choosing forwards results in an important simplification

Ways to measure variance (3): simple variance swaps

Result (Pricing, version 1)

The strike on a simple variance swap is

$$V = \sum_{i=1}^{T/\Delta} \frac{e^{ri\Delta}}{F^2_{(i-1)\Delta}} \left[\Pi(i\Delta) - (2 - e^{-r\Delta}) \Pi((i-1)\Delta) \right] \tag{1}$$

where $\Pi(t)$ *is given by*

$$\Pi(t) = 2 \int_0^{F_t} \mathrm{put}_t(K)\, dK + 2 \int_{F_t}^{\infty} \mathrm{call}_t(K)\, dK + S_0^2 e^{rt}$$

Ways to measure variance (3): simple variance swaps

- This result does not require the price process of the underlying asset to be continuous
- But the hedging portfolio requires holding portfolios of options of each maturity: $\Delta, 2\Delta, \ldots, T-\Delta, T$
- Not a serious issue if Δ is large relative to T, but raises the concern that hedging a simple variance swap may be extremely costly in practice if Δ is very small relative to $T \ldots$?

Ways to measure variance (3): simple variance swaps

Result (Pricing, version 2)

In the limit as $\Delta \to 0$,

$$V = \frac{2e^{-rT}}{S_0^2} \left\{ \int_0^{F_T} \text{put}_T(K)\, dK + \int_{F_T}^{\infty} \text{call}_T(K)\, dK \right\}$$

- Hedge:
  - (i) a static position in options expiring at T: hold OTM puts and OTM calls, equally weighted
  - (ii) a dynamic delta-hedge
- Hedging and pricing works even if the underlying asset's price can jump

A sample path

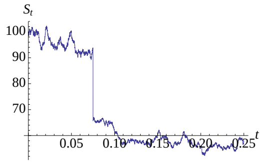

Figure: SVS: price process

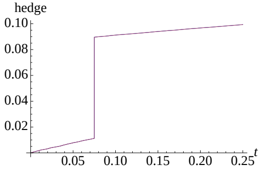

Figure: Hedge vs. payoff

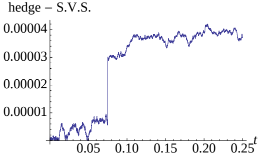

Figure: Hedge error

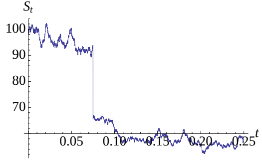

Figure: VS: price process

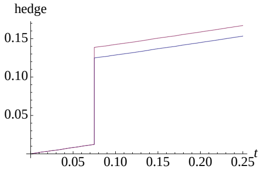

Figure: Hedge vs. payoff

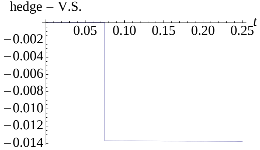

Figure: Hedge error

Robustness of simple variance swaps

What if $\Delta > 0$?

- In the limit as $\Delta \to 0$, we saw that the fair strike on a simple variance swap is

$$V = \frac{2e^{-rT}}{S_0^2} \left\{ \int_0^{F_T} \text{put}_T(K)\, dK + \int_{F_T}^{\infty} \text{call}_T(K)\, dK \right\}$$

- In practice, contract has to stipulate $\Delta > 0$
- Paper derives analytic bound showing that the error in approximating $V(\Delta)$ by the above V is very small (even if there are jumps)
- Percentage error $\sim 0.008\%$ with daily sampling interval, $\Delta = 1/252$

Robustness of simple variance swaps

What if you can't trade deep out-of-the-money strikes?

- What if you can only trade strikes in the range (A, B)?
- Modify payoff with a correction term

$$\left(\frac{S_\Delta - S_0}{S_0}\right)^2 + \left(\frac{S_{2\Delta} - S_\Delta}{F_\Delta}\right)^2 + \cdots + \left(\frac{S_T - S_{T-\Delta}}{F_{T-\Delta}}\right)^2 - \phi(S_T)$$

where

$$\phi(S_T) = \begin{cases} \left(\frac{A - S_T}{F_{T-\Delta}}\right)^2 & \text{if } S_T < A \\ 0 & \text{if } A \le S_T \le B \\ \left(\frac{S_T - B}{F_{T-\Delta}}\right)^2 & \text{if } S_T > B \end{cases}$$

- Can be replicated with strikes in the range (A, B)
- For 1-month simple variance swaps, correction term was zero on every date in my sample

Robustness of simple variance swaps

What if you can't trade deep out-of-the-money strikes?

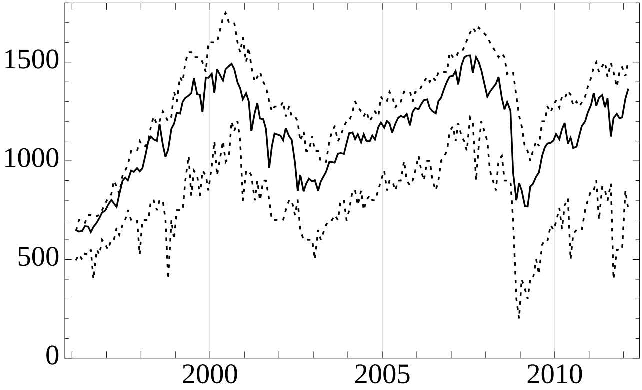

Figure: Dashed lines indicate strike range tradable on a given day. Solid line indicates where the market ended up 30 days later

What does VIX measure?

- The VIX index is based on the strike on a variance swap:

$$\text{VIX}^2 = \frac{2e^{rT}}{T}\left\{\int_0^{F_T} \frac{1}{K^2}\,\text{put}_T(K)\,dK + \int_{F_T}^{\infty} \frac{1}{K^2}\,\text{call}_T(K)\,dK\right\}$$

- This is a *definition*, not a statement about pricing
- Generally interpreted as a measure of risk-neutral variance $\mathbb{E}^* \int_0^T \sigma_t^2\,dt$
- In the presence of jumps...
  - ...the replicating portfolio doesn't replicate the variance swap payoff
  - ...the correctly priced strike $\widetilde{V}$ doesn't equal VIX^2
  - ...neither $\widetilde{V}$ nor VIX^2 has the interpretation $\mathbb{E}^* \int_0^T \sigma_t^2\,dt$

What does VIX measure?

Result (Interpretation of VIX)

VIX measures the entropy of the simple return $R_T = S_T/S_0$*,*

$$VIX^2 = \frac{2}{T}\left(\log \mathbb{E}^* R_T - \mathbb{E}^* \log R_T\right)$$

If R_T *is lognormal, then*

$$VIX^2 = \frac{1}{T}\,\mathrm{var}^* \log R_T \approx \frac{1}{T}\,\mathrm{var}^* R_T$$

But, in general, with jumps and/or time-varying volatility, VIX depends on all of the (annualized, risk-neutral) cumulants of log returns,

$$VIX^2 = 2\sum_{n=2}^{\infty} \frac{\kappa_n^*}{n!} = \kappa_2^* + \frac{\kappa_3^*}{3} + \frac{\kappa_4^*}{12} + \frac{\kappa_5^*}{60} + \cdots$$

What does VIX measure?

- Initially surprising observation: negative skewness drives VIX down

$$\text{VIX}^2 = 2\sum_{n=2}^{\infty} \frac{\kappa_n^*}{n!} = \kappa_2^* + \frac{\kappa_3^*}{3} + \frac{\kappa_4^*}{12} + \frac{\kappa_5^*}{60} + \cdots$$

What does VIX measure?

- Initially surprising observation: negative skewness drives VIX down

$$\text{VIX}^2 = 2\sum_{n=2}^{\infty} \frac{\kappa_n^*}{n!} = \kappa_2^* + \frac{\kappa_3^*}{3} + \frac{\kappa_4^*}{12} + \frac{\kappa_5^*}{60} + \cdots$$

- But this is skewness *calculated with risk-neutral probabilities*
- To see how VIX is linked to real-world probabilities, suppose that there is a representative investor with log utility

What does VIX measure?

Result (VIX in equilibrium)

If there is a representative agent with log utility, then VIX can be expressed in terms of the cumulants of $\log R_T$ *under the real-world probabilities,*

$$VIX^2 = 2\sum_{n=2}^{\infty}(-1)^n(n-1)\frac{\kappa_n}{n!} = \kappa_2 - \frac{2}{3}\kappa_3 + \frac{\kappa_4}{4} - \frac{\kappa_5}{15} + \cdots$$

What does SVIX measure?

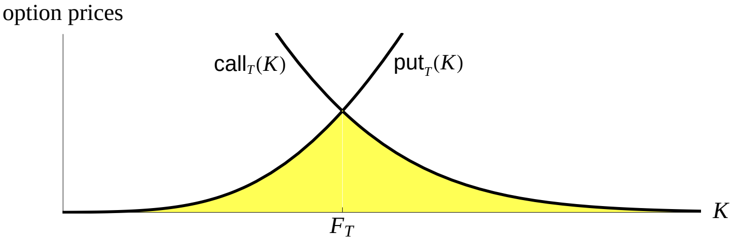

- Analogously, we can define an index, SVIX, based on the strike of a simple variance swap:

$$\text{SVIX}^2 = \frac{2e^{rT}}{T \cdot S_0^2} \left\{ \int_0^{F_T} \text{put}_T(K)\, dK + \int_{F_T}^{\infty} \text{call}_T(K)\, dK \right\}$$

What does SVIX measure?

Result (Interpretation of SVIX)

SVIX measures the risk-neutral variance of the simple return:

$$SVIX^2 = \frac{1}{T} \operatorname{var}^* R_T$$

VIX and SVIX

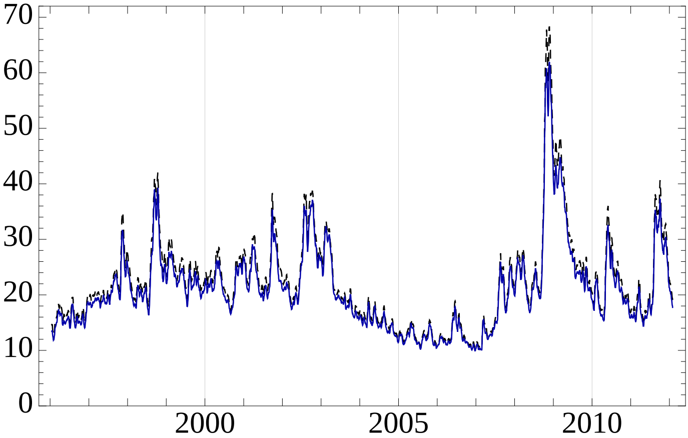

Figure: VIX (dotted) and SVIX (solid). Jan 4, 1996–Jan 31, 2012
Figure shows 10-day moving average. $T = 1$ month

VIX minus SVIX

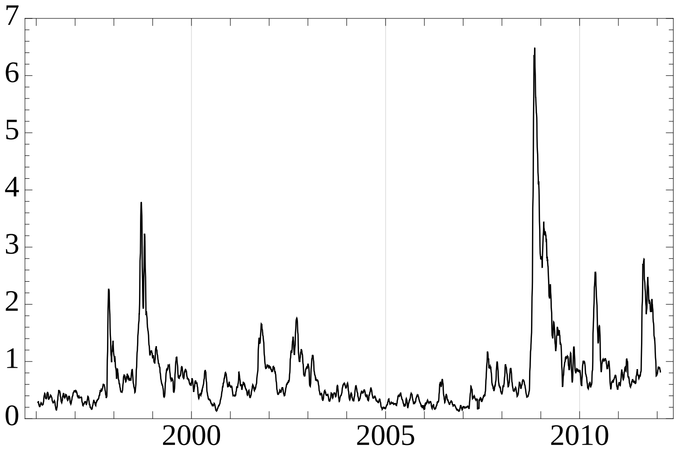

Figure: VIX minus SVIX. Jan 4, 1996–Jan 31, 2012

Figure shows 10-day moving average. $T = 1$ month

VIX minus SVIX

- If returns and the SDF are conditionally lognormal with return volatility σ_R then we can calculate VIX and SVIX in closed form:

$$\begin{aligned} \text{SVIX}^2 &= \frac{1}{T} e^{2rT} \left(e^{\sigma_R^2 T} - 1 \right) \\ \text{VIX}^2 &= \sigma_R^2 \end{aligned}$$

- VIX would be lower than SVIX
- *No* conditionally lognormal model can match the data

A lower bound on the equity premium

- SVIX gives a bound on the equity premium perceived by an investor
  - who is unconstrained,
  - who holds the market over the horizon of interest, and
  - whose relative risk aversion (which need not be constant) is at least one
- I have nothing to say about
  - Constrained investors
  - Irrational investors
  - Investors who don't hold the market
  - The connection between prices and cashflows or consumption

A lower bound on the equity premium

- Notation: $\frac{1}{R_{f,T}} \mathbb{E}^* X_T = \mathbb{E}\, M_T X_T$
- Start from the following identity:

$$
\begin{aligned}
\frac{\operatorname{var}^* R_T}{R_{f,T}} &= \frac{1}{R_{f,T}} \mathbb{E}^* R_T^2 - \frac{1}{R_{f,T}} \left(\mathbb{E}^* R_T\right)^2 \\
&= \mathbb{E}(M_T R_T^2) - R_{f,T} \\
&= \mathbb{E}\, R_T - R_{f,T} + \mathbb{E}(M_T R_T^2) - \mathbb{E}\, R_T \\
&= \mathbb{E}\, R_T - R_{f,T} + \underbrace{\operatorname{cov}(M_T R_T, R_T)}_{\leq 0}
\end{aligned}
$$

- This connects something we can measure to something interesting plus something we can control

A lower bound on the equity premium

Result (A bound on the equity premium, version 1)

If there is a one-period investor who holds the market from time 0 to time T, and whose risk aversion $\gamma(c) \equiv -cu''(c)/u'(c) \geq 1$*, then*

$$\mathbb{E}\left[R_T - R_{f,T}\right] \geq \frac{\text{var}^* R_T}{R_{f,T}}$$

With log utility, $\gamma(x) \equiv 1$*, this holds with equality*

- Equivalently, $\frac{1}{T}\,\mathbb{E}\left[R_T - R_{f,T}\right] \geq \frac{\text{SVIX}^2}{R_{f,T}}$
- Does not require power utility: γ doesn't have to be constant
- Does not require lognormality

A lower bound on the equity premium

Proof.

- The given assumption implies that the SDF is proportional to $u'(R_T)$, so we must show that $\text{cov}\left(R_T u'(R_T), R_T\right) \leq 0$
- This holds because $R_T u'(R_T)$ is decreasing in R_T: its derivative is $u'(R_T) + R_T u''(R_T) = -u'(R_T)\left[\gamma(R_T) - 1\right] \leq 0$

□

A lower bound on the equity premium

- Paper extends to the case of an intertemporal consumer-investor:

$$V[W] = \max_{C,w_i} u(C) + \beta \, \mathbb{E} \, V\left[(W - C) \sum_i w_i R_T^{(i)}\right]$$

- The right coefficient of risk aversion is not $-\frac{Cu''(C)}{u'(C)}$ but $-\frac{WV''(W)}{V'(W)}$
- The envelope condition, $u'(C) = V'(W)$, implies that

$$\underbrace{-\frac{WV''(W)}{V'(W)}}_{\text{must be} \geq 1} = -\frac{Cu''(C)}{u'(C)} \times \frac{\partial \log C}{\partial \log W}$$

A lower bound on the equity premium

- If your preferred model is Bansal–Yaron or Campbell–Cochrane, the above does not apply: it requires separable utility
- But we can (approximately) deal with these cases too: under conditional lognormality,

$$\frac{1}{T}\,\mathbb{E}\left[R_T - R_{f,T}\right] \geq \frac{\text{SVIX}^2}{R_{f,T}}$$

  if and only if the conditional Sharpe ratio is greater than the conditional volatility of the market—which holds in the data
- The inequality also holds with Epstein–Zin preferences whenever $\gamma \geq 1$ and EIS is sufficiently close to 1 (no need for lognormality)

A lower bound on the equity premium

$$\frac{\text{var}^* R_T}{R_{f,T}} \le \mathbb{E}\, R_T - R_{f,T} \le R_{f,T} \cdot \sigma(M) \cdot \sigma(R_T)$$

- Left-hand inequality is the new result
  - Good: relates unobservable equity premium to an observable quantity
  - Bad: requires an economic assumption
- Right-hand inequality is the Hansen–Jagannathan bound
  - Good: no assumptions
  - Bad: neither side is observable

A lower bound on the equity premium

- Merton (1980) estimated equity premium from

$$\text{instantaneous risk premium} = \gamma\sigma^2$$

- Assumes power utility and the market's price follows a diffusion
- No distinction between risk-neutral and real-world variance in a diffusion-based model (Girsanov's theorem)
- Beyond diffusions, the appropriate generalization relates the risk premium to the *risk-neutral* variance

A lower bound on the equity premium

1mo horizon, annualized, 10-day moving avg. Mid prices in black, bid prices in red

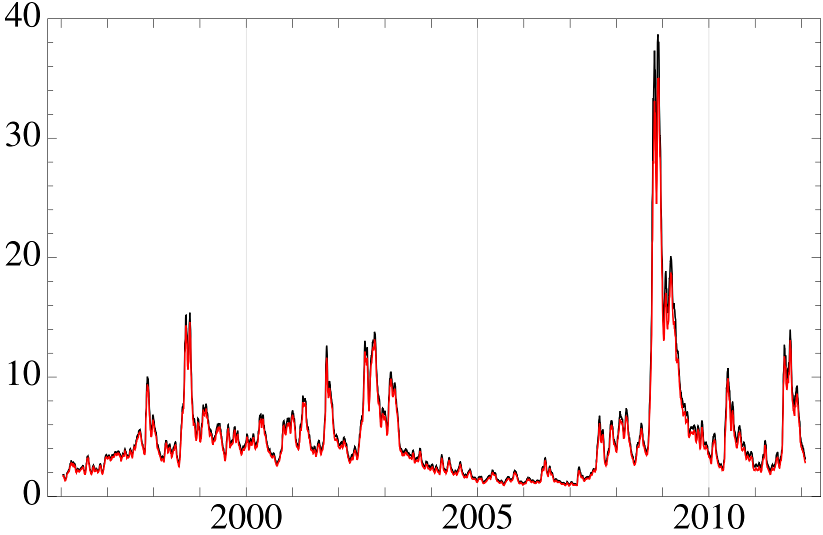

A lower bound on the equity premium

3mo horizon, annualized, 10-day moving avg. Mid prices in black, bid prices in red

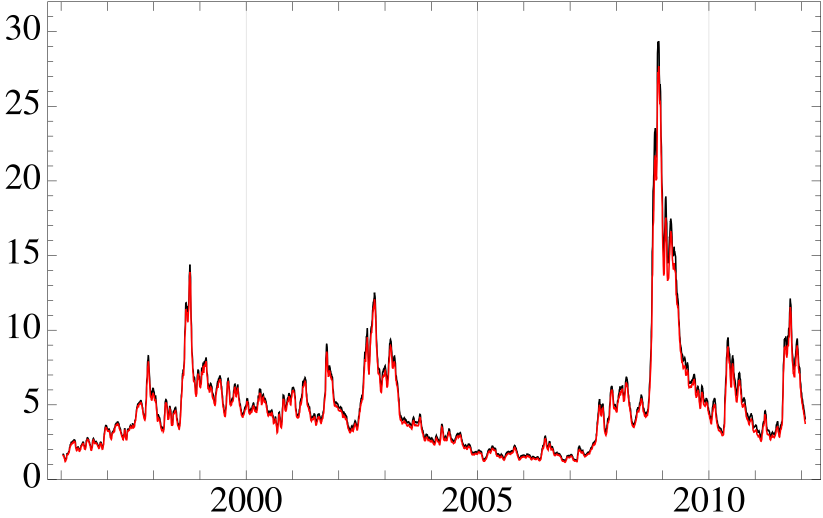

A lower bound on the equity premium

1yr horizon, annualized, 10-day moving avg. Mid prices in black, bid prices in red

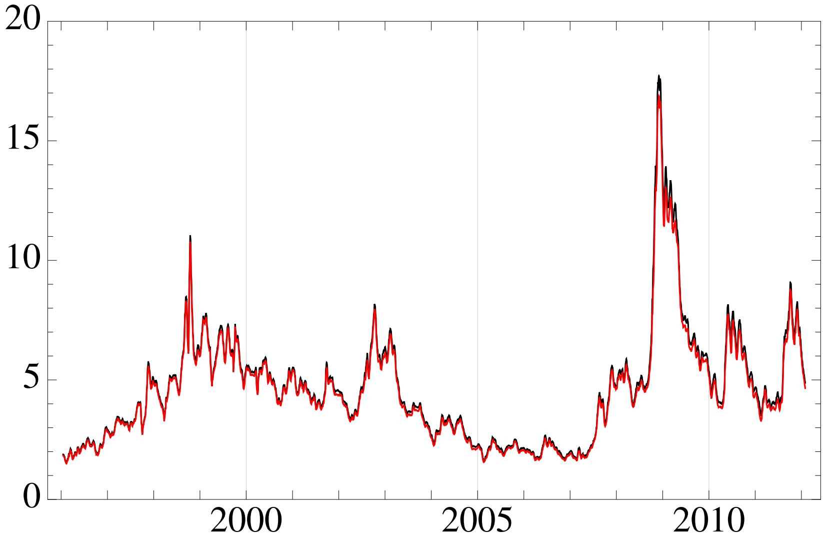

A lower bound on the equity premium

horizon	mean	s.d.	min	1%	10%	25%	50%	75%	90%	99%	max
1 mo	4.92	4.53	0.81	1.01	1.52	2.41	3.85	5.65	8.85	25.4	54.2
2 mo	4.92	3.94	0.97	1.17	1.62	2.57	4.05	5.82	8.41	23.2	45.4
3 mo	4.87	3.55	1.03	1.25	1.70	2.65	4.15	5.83	8.05	21.1	38.6
6 mo	4.77	2.93	1.24	1.45	1.88	2.83	4.29	5.88	7.57	16.7	28.6
1 yr	4.48	2.41	1.03	1.53	1.95	2.71	4.19	5.48	7.02	13.7	21.2

Table: Mean, standard deviation, and quantiles of EP bound (in %)

A lower bound on the equity premium

- Equity premium was very high at times of stress
- Also high from 1998–2000
  - Forecasts based on market valuation ratios incorrectly predicted a low or even negative equity premium during this period
  - By construction, the lower bound can never be less than zero
- Most important: no out-of-sample issues (Ang–Bekaert, Goyal–Welch) because no parameter estimation is required

A lower bound on the equity premium

- Suppose you think this just reflects "market segmentation". What *trade* should you have done in November 2008?
  - Short the portfolio of options, i.e. short an at-the-money-forward straddle and (equally weighted) out-of-the-money calls and puts
  - You end up short if the market rallies and long if the market sells off
  - You're taking a contrarian position, providing liquidity to the market
- At the height of the credit crisis, extraordinarily high risk premia were available for investors who were prepared to take on this position

How *not* to think about what's going on

- If the world were lognormal you could think of what we're doing as
  - Exploiting a relationship between risk premia and $\mathbb{P}$-variance
  - $\mathbb{Q}$-variance equals $\mathbb{P}$-variance
- In the real, non-lognormal, world, this is the wrong intuition
  - Risk premia sensitive to higher moments as well as $\mathbb{P}$-variance
  - $\mathbb{Q}$-variance not equal to $\mathbb{P}$-variance

The bound is conservative in two respects

- Can't observe deep-OTM option prices

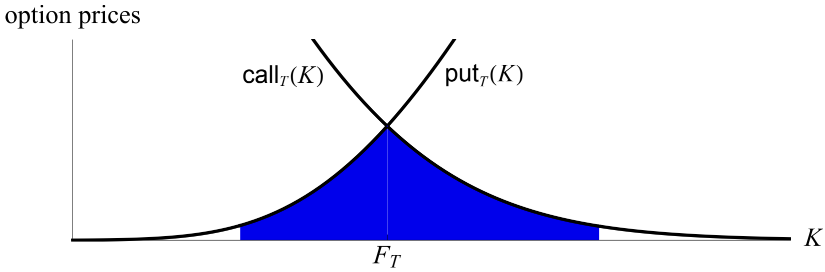

The bound is conservative in two respects

- Even near-the-money, can't observe a continuum of strikes

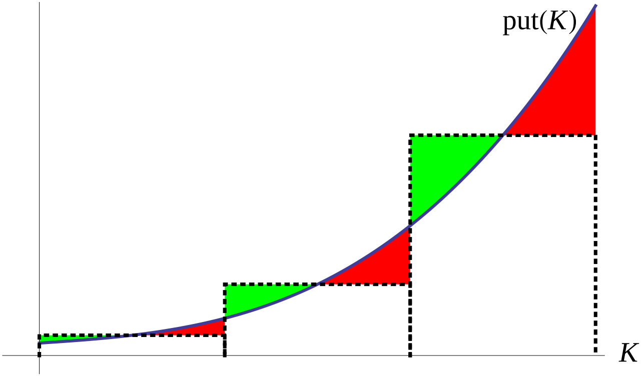

The bound is conservative in two respects

- Both these effects mean that the true lower bound is even higher
- By ignoring deep-OTM options, we underestimate the true area under the curve
- More subtly, discretization in strike also leads to underestimating the true area, because call(K) and put(K) are both convex in K

Labour income

- Suppose the investor has some other source of income $L(R_T)$ at time T; think of bonds or labour income
- Then we get the same result under a stronger assumption on risk aversion
  - If $L'(R_T) \geq 0$ and $L(R_T) \leq \kappa W_T$ then we need risk aversion at least $1 + \kappa$
  - If the agent has at least as much wealth in the market as labour (or bond) income between now and time T, $L(R_T) \leq W_T$, then we need risk aversion at least 2
  - All the results go through the same—and the *numbers* are the same too

Equity premium bound (rescaled) and S&P 500

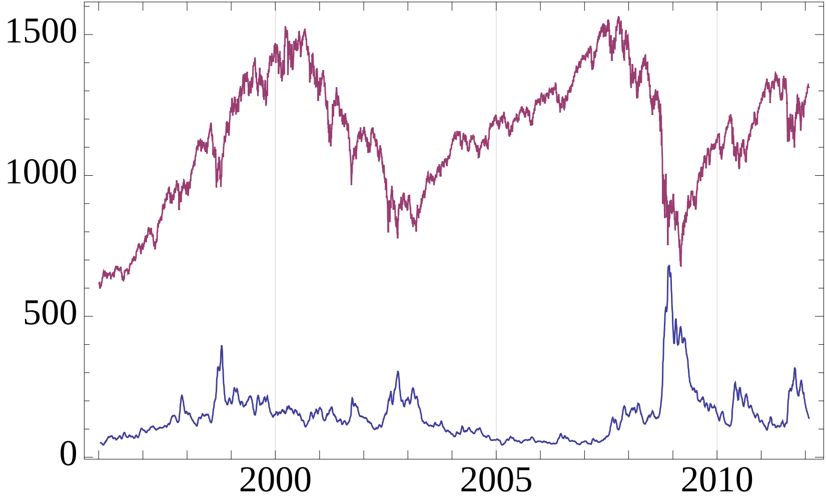

What if the bound were tight?

- $REP_{t \to t+T} \overset{?}{\approx} \alpha + \beta \cdot EPB_t + \varepsilon_{t+T}$
- OLS with Hansen–Hodrick standard errors to account for heteroskedasticity and overlapping observations
- Null hypothesis: bound holds with equality, $\alpha = 0, \beta = 1$

horizon	$\widehat{\alpha}$	s.e.	$\widehat{\beta}$	s.e.	R^2
1 mo	-0.01	[0.06]	0.77	[1.41]	0.3%
2 mo	-0.02	[0.07]	0.95	[1.47]	0.8%
3 mo	-0.02	[0.07]	0.90	[1.63]	0.9%
6 mo	-0.07	[0.06]	1.90	[0.90]	4.6%
1 yr	-0.04	[0.09]	1.54	[1.29]	3.6%

Up-SVIX and down-SVIX

Work in progress

$$\operatorname{var}^* R_T = \underbrace{\frac{2R_{f,T}}{P_0^2} \int_0^{F_{0,T}} \text{put}_T(K)\, dK}_{\text{“down-SVIX}^2\text{”}} + \underbrace{\frac{2R_{f,T}}{P_0^2} \int_{F_{0,T}}^{\infty} \text{call}_T(K)\, dK}_{\text{“up-SVIX}^2\text{”}}$$

- Natural to think about what the calls and puts tell us separately
- Can we give a nice interpretation to down- and up-SVIX^2?

Up-SVIX and down-SVIX

- Yes, if we think from the perspective of a log investor

$$
\begin{aligned}
\mathbb{P}\left(R_T > R_{f,T}\right) &= -R_{f,T}\,\text{call}'(F_{0,T}) + \frac{\text{call}(F_{0,T})}{P_0} \\
\mathbb{E}\left[(R_T - R_{f,T})\mathbf{1}\left\{R_T > R_{f,T}\right\}\right] &= \frac{R_{f,T}}{P_0}\,\text{call}(F_{0,T}) + \frac{\text{up-SVIX}^2}{R_{f,T}} \\
\mathbb{E}\left[(R_T - R_{f,T})\mathbf{1}\left\{R_T < R_{f,T}\right\}\right] &= -\frac{R_{f,T}}{P_0}\,\text{call}(F_{0,T}) + \frac{\text{down-SVIX}^2}{R_{f,T}}
\end{aligned}
$$

- These are real-world probabilities, $\mathbb{P}$ not $\mathbb{P}^*$
- Rule of thumb: $\mathbb{E}\left[R_T - R_{f,T} \middle| R_T > R_{f,T}\right] \approx 2 \times \text{up-SVIX}^2$
- $\mathbb{E}\left[R_T - R_{f,T} \middle| R_T < R_{f,T}\right] \approx 2 \times \text{down-SVIX}^2$

Probability of an up-move, $\mathbb{P}\left(R_T > R_{f,T}\right)$

$T = 1$ mo

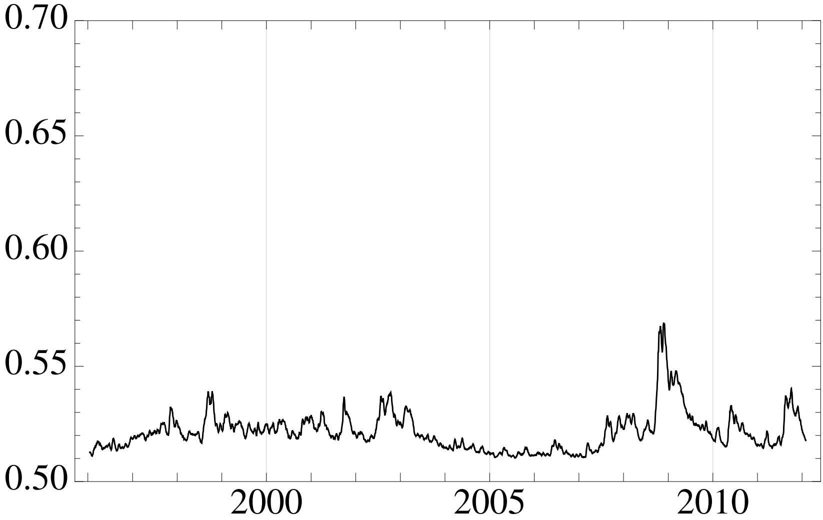

Probability of an up-move, $\mathbb{P}\left(R_T > R_{f,T}\right)$

$T = 1$ yr

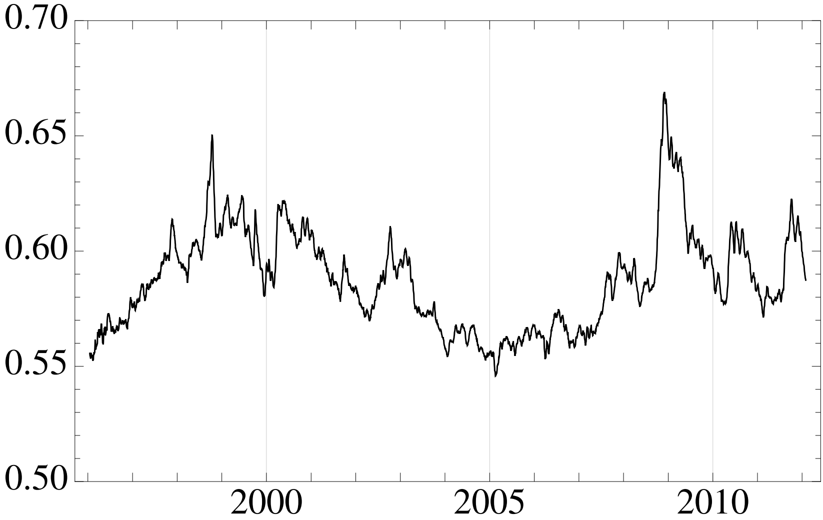

Conditional equity premium $\mathbb{E}\left(R_T - R_{f,T} \middle| R_T \gtrless R_{f,T}\right)$

Black: up-move. Red: down-move (sign flipped). $T = 1$ yr

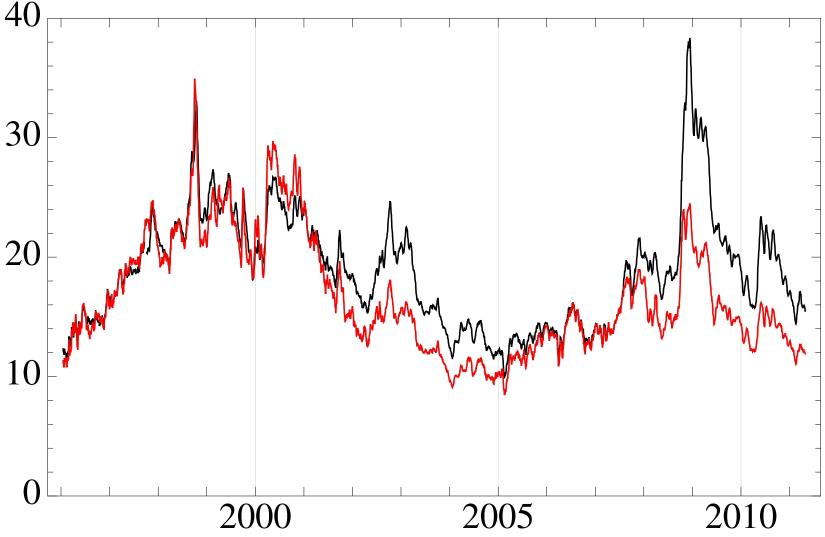

A comparison: November 1998 vs November 2008

- Both were times at which the equity premium was high
- In both cases, the probability of an up-move was high (over 65%)
- Expected excess return conditional on an up-move was high in both cases (33% in ’98, 38% in ’08)
- Big gap in the expected excess return conditional on a down-move (-35% in ’98, -24% in ’08)
- Internet boom was the only period in the sample when (the absolute value of) the equity premium was larger conditional on a down-move than on an up-move

Generalizing the bound

Work in progress

Modify the previous assumptions as follows:

1. There is a one-period investor as before, except that relative risk aversion $\gamma(x) \equiv -xu''(x)/u'(x)$ satisfies $\gamma(x) \in [\underline{\gamma}, \overline{\gamma}]$, or
2. There is an intertemporal investor as before, except that relative risk aversion $\Gamma(x) \equiv -xJ''[x]/J'[x]$ satisfies $\Gamma(x) \in [\underline{\gamma}, \overline{\gamma}]$, or
3. There is an Epstein-Zin (1989) investor as before, except that risk aversion $\gamma \in [\underline{\gamma}, \overline{\gamma}]$, or
4. The SDF and market return are conditionally jointly lognormal, and the ratio of the market's conditional Sharpe ratio to its conditional volatility lies between $\underline{\gamma}$ and $\overline{\gamma}$

Generalizing the bound

Result

Under any one of these assumptions,

$$\frac{\mathbb{E}^*\left(R_T^{\theta+\underline{\gamma}}\right)}{\mathbb{E}^*\left(R_T^{\underline{\gamma}}\right)} \leq \mathbb{E}\left(R_T^{\theta}\right) \leq \frac{\mathbb{E}^*\left(R_T^{\theta+\overline{\gamma}}\right)}{\mathbb{E}^*\left(R_T^{\overline{\gamma}}\right)} \qquad \textit{for any } \theta \geq 0$$

and

$$\frac{\mathbb{E}^*\left(R_T^{\theta+\overline{\gamma}}\right)}{\mathbb{E}^*\left(R_T^{\overline{\gamma}}\right)} \leq \mathbb{E}\left(R_T^{\theta}\right) \leq \frac{\mathbb{E}^*\left(R_T^{\theta+\underline{\gamma}}\right)}{\mathbb{E}^*\left(R_T^{\underline{\gamma}}\right)} \qquad \textit{for any } \theta \leq 0.$$

If one of the bounds on relative risk aversion holds with equality, then the corresponding inequality holds with equality.

Bounds on the equity premium

Bounds on 1yr equity premium, 10-day moving average, $\gamma = 1, 2, 4$

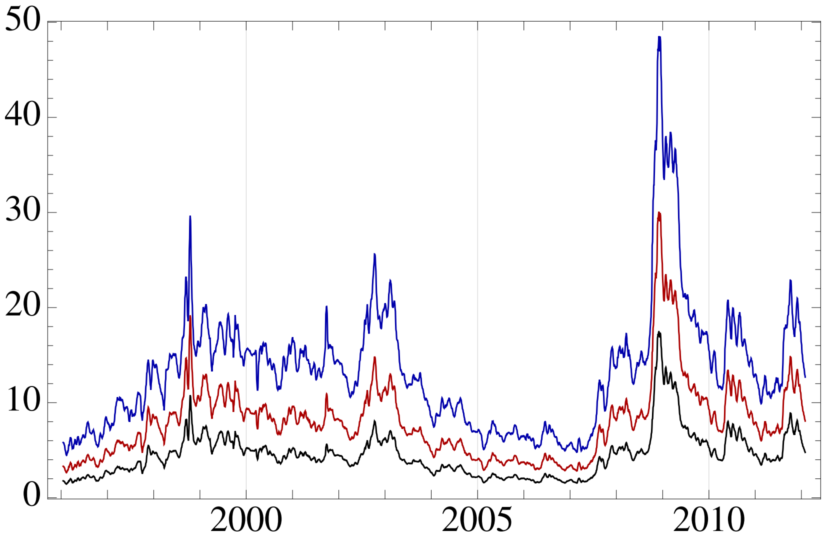

Bounds on the equity premium

γ	mean	s.d.	min	1%	10%	25%	50%	75%	90%	99%	max
0.75	3.47	1.89	0.75	1.17	1.50	2.09	3.23	4.25	5.49	10.7	16.6
1	4.48	2.41	1.03	1.53	1.95	2.71	4.19	5.48	7.02	13.7	21.2
1.25	5.44	2.89	1.33	1.87	2.39	3.30	5.10	6.65	8.44	16.5	25.5
1.50	6.34	3.34	1.65	2.20	2.80	3.86	5.98	7.76	9.74	19.0	29.4
2	8.01	4.15	2.34	2.81	3.58	4.90	7.60	9.79	12.2	23.7	36.5
4	13.4	6.71	3.79	4.85	6.16	8.24	12.9	16.2	19.9	38.7	59.3
8	21.1	11.4	4.47	7.27	9.70	12.7	20.5	24.6	31.7	70.7	114.8

Table: Mean, standard deviation, and quantiles of equity premium bounds at the 1-year horizon (measured in %) for different levels of risk aversion

Conclusion

- Proposed a new derivative contract, the simple variance swap, that can be hedged and therefore priced even if there are jumps
- Constructed an index, SVIX, that is to a simple variance swap as VIX is to a variance swap
- SVIX is less than VIX in the data; conditionally lognormal models make the opposite prediction
- No need to assume that the world is stationary or ergodic, because there's no need to replace $\mathbb{E}$ with $\frac{1}{T}\sum$
- The forward-looking equity premium is extremely volatile and was extraordinarily high during the crisis