The (Time-Varying) Importance of Disaster Risk

Frankfurt Goethe Seminar

Ivo Welch

June 2015

Prominent Research Area

Dark Events, Catastrophes, Disasters, Tail-Risk, Black Swans

- Barro 2006 (intl disasters)
- Broadie Chernov Johannes 2009 (price kernels)
- Chen Dou Kogan 2013
- Gao Song 2013 (exposure of firms to OOM options)
- Gourio 2012 (business cycles)
- Kelley and Jiang 2015 (firms loading on tail risk)
- (Kozeniauskas,) Orlik, Verldkamp (model time-variation)
- Rietz 1988 (original suggestions)
- Taleb 2001 (prominent book)
- Wachter 2013
- ....

Warning

- As always, experts probably are already **in the know**.
- I am not an expert. I just saw too many seminars about disasters.
- I am sticking my nose into something that is not my expertise
- I am pretty slow thinking about this stuff.

Simple

When you can buy insurance against disasters of a particular kind at a cost of 1-2% per annum (of the value of your portfolio), then it cannot be that fear of these disasters can explain more than 1-2% (of the return on your portfolio).

My paper focuses on post-1983 data, in which the geometric equity (net of rf) premium was 7.2% per annum.

Recent Index Put Option Prices

- April 20, 2015. S&P 500 = 2086.20.
- Cost of a Put With Strike 1500 (–22% OOM), 85-day maturity:

  $0.70

  (A 22% loss would reduce long-term geo-mean by $< 0.1\%$/yr.)
- 4.3 (365/85) times $0.70 is **$3/year** on 2086.20 index.
- $3/2086 \approx 0.15\%$.
- A full 1-year protection at –15% is $< 1\%$/year. Protects against multiple, say, –90% crashes.
- How likely can –90% disasters be *today*? How much will the market pay you in the up-states to take this kind of risk in the down-states?
- This is most of what this paper has to say.

Recent Index Put Option Prices

- April 20, 2015. S&P 500 = 2086.20.
- Cost of a Put With Strike 1500 (–22% OOM), 85-day maturity:

  $0.70

  (A 22% loss would reduce long-term geo-mean by $<$ 0.1%/yr.)
- 4.3 (365/85) times $0.70 is **$3/year** on 2086.20 index.
- $3/2086 \approx 0.15\%$.
- A full 1-year protection at –15% is $<$ 1%/year. Protects against multiple, say, –90% crashes.
- How likely can –90% disasters be *today*? How much will the market pay you in the up-states to take this kind of risk in the down-states?
- This is most of what this paper has to say.

What Does It Protect Against?

Like earthquake insurance, not perfect:

- No protection against repeated –20%/year drops.
  (But future fears should be immediately incorporate into stock prices. We have serial correlation in volatility, not in returns.)
- No protection against higher renewal prices. No protection against time-varying changes in long-run expected returns.
- No protection against default of protection seller.

This far-below-the-money put-based stock protection is very specifically against an intra-monthly large stock market crash.

What Does It Protect Against?

Like earthquake insurance, not perfect:

- No protection against repeated –20%/year drops.
  (But future fears should be immediately incorporate into stock prices. We have serial correlation in volatility, not in returns.)

- No protection against higher renewal prices. No protection against time-varying changes in long-run expected returns.

- No protection against default of protection seller.

This far-below-the-money put-based stock protection is very specifically against an intra-monthly large stock market crash.

What Does It Protect Against?

Like earthquake insurance, not perfect:

- No protection against repeated –20%/year drops.
  (But future fears should be immediately incorporate into stock prices. We have serial correlation in volatility, not in returns.)
- No protection against higher renewal prices. No protection against time-varying changes in long-run expected returns.
- No protection against default of protection seller.

This far-below-the-money put-based stock protection is very specifically against an intra-monthly large stock market crash.

What Does It Protect Against?

Like earthquake insurance, not perfect:

- No protection against repeated –20%/year drops.
  (But future fears should be immediately incorporate into stock prices. We have serial correlation in volatility, not in returns.)
- No protection against higher renewal prices. No protection against time-varying changes in long-run expected returns.
- No protection against default of protection seller.

This far-below-the-money put-based stock protection is very specifically against an intra-monthly large stock market crash.

Historical Data Option

Far-Below-The-Money Index Put Prices

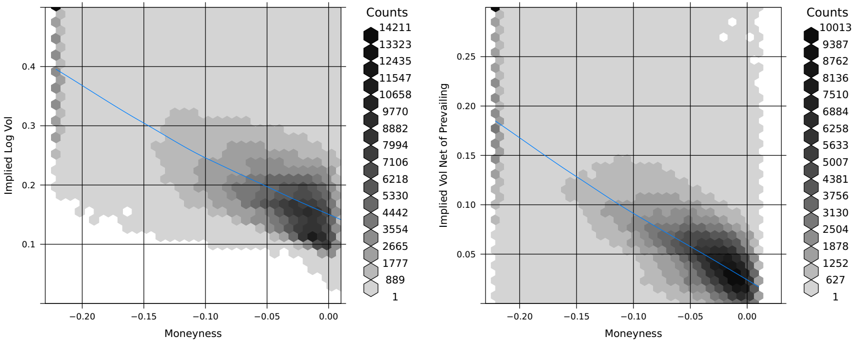

Left Side of Volatility Smirk. 15% BTM options (30-day) were priced at a vol of 30%, much more than the ATM vol of 15%.

Additional Net Tail Risk in Crisis?

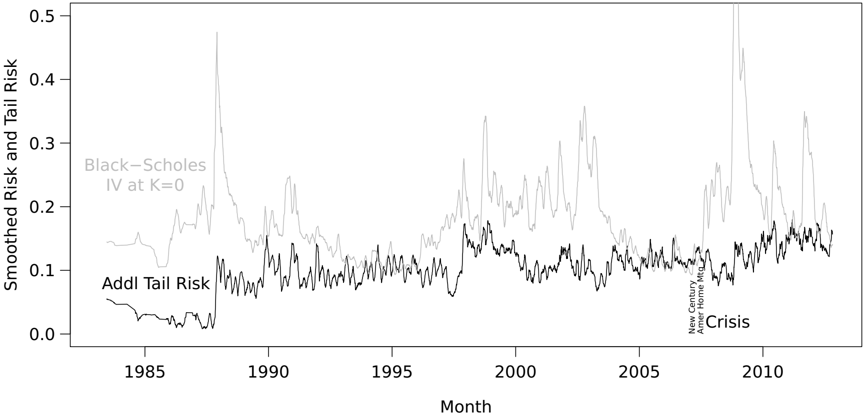

Left-tail risk did not jump up more than general risk during the financial crisis.

Prices or Implied Vol?

- I care about the (absolute dollar) price of put protection.
- I care less directly about the implied volatility.
- A deeper BTM put is better for insuring against a –100% crash, because it is still cheaper in dollar terms, even if it is more expensive in Black-Scholes terms.

Since 1983

- 1983-2013 (320 mos), geo eq-prem: 0.58%/mo ($\approx$ 7.1%/yr).
  [1926-2014 = 0.51%; 2004-2014 = 0.56%]
- Careful CME 15%-BTM data, 1983-2013, 1-mo put, 15% BTM.
- Annual Total Historical Cost, Always Buying Insurance:

$$< 1.5\%/\text{yr}.$$

- Disaster risk was important (well, protection was expensive), but it cannot explain a 7%/year equity premium.
  Crash Risk was probably less important than ordinary sampling variation (s.e.: 2%/year).
- When disaster insurance was too expensive, an investor could have gotten out. Yes, some good returns would have been lost, but the included periods were still offering about 6% per annum.

Time Variation in Disaster Premia

- The time-varying price of 1-mo –15% disaster insurance:

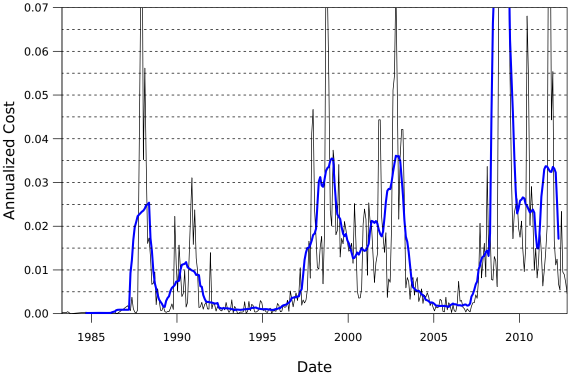

- Disaster risk is very much time-varying. Insurance can be > 2%/yr.
- When high, you could have gotten out of stocks+put.

Time-Varying Investment Strategy

- Maybe not a great idea to buy disaster insurance if P is high?
- If you just get out of the stock market in high-insurance-cost months, cost of insurance goes from $\approx$ 1.5%/yr to $\approx$ 1% year.
- but would also have been out of the market in good months:
- Both: 7.1% uncond. geo eqprem vs. 5.8% put-protected

  Clearly, you would have sacrificed some returns being risk averse. But net cost is now 1.3%.

- You would still have earned a geometric 4.5%/year.
- Limits disaster risk a little more, but not greatly.
- By definition, realized sd and mean are unfortunately not expected, so I cannot tell you how much the put reduced your risk. I know it was pricy.

  (I cannot look at strategies in which I stay in the market, but do not protect if the put is too expensive. Ex-post realization was not representative.)

What I am sparing you

- Details on how to calculate and smooth below-the-money put option prices.
- Not frequently traded. Need to be aligned **intra-**day.

I will post the data soon.

What Can Be Said About Magnitude and Frequency?

- A time-varying risk-neutral price consists of a time-varying "fear" component and a time-varying "frequency" component.
- I need a prior model for the frequency of disaster or a good equilibrium pricing model.

- I don't have a good equilibrium pricing model.
- No one has one, which is a bit ironic given that the models usually presume we know it.

- So, I will assume a prior model for the frequency of disasters.

What Can Be Said About Magnitude and Frequency?

- A time-varying risk-neutral price consists of a time-varying "fear" component and a time-varying "frequency" component.
- I need a prior model for the frequency of disaster or a good equilibrium pricing model.

- I don't have a good equilibrium pricing model.
- No one has one, which is a bit ironic given that the models usually presume we know it.

- So, I will assume a prior model for the frequency of disasters.

What Can Be Said About Magnitude and Frequency?

- A time-varying risk-neutral price consists of a time-varying "fear" component and a time-varying "frequency" component.
- I need a prior model for the frequency of disaster or a good equilibrium pricing model.

- I don't have a good equilibrium pricing model.
- No one has one, which is a bit ironic given that the models usually presume we know it.

- So, I will assume a prior model for the frequency of disasters.

Prior Probability Choices

- Dogmatic — the frequency (of losing all your money) is 5% (or 0.1% or 99.9%). No data can reject your prior.
- Parametric — (e.g., Veldkamp). Probability is lognormal, and we can update and compute the left-tail, even when we draw right-tail observations.
- Diffuse — In 1980, you had a probability belief of
  - 1/323 ≈ 0.3% that the per-month disaster risk was 0.0%.
  - 1/323 ≈ 0.3% that the per-month disaster risk was 0.3%.
  - 1/323 ≈ 0.3% that the per-month disaster risk was 0.6%.
  - 1/323 ≈ 0.3% that the per-month disaster risk was 0.9%.
  - ...
  - 1/323 ≈ 0.3% that the per-month disaster risk was 100.0%.
- All regression results assume similar diffuse priors. With dogmatic priors, regressions are not informative.

Prior Probability Choices

- Dogmatic — the frequency (of losing all your money) is 5% (or 0.1% or 99.9%). No data can reject your prior.
- Parametric — (e.g., Veldkamp). Probability is lognormal, and we can update and compute the left-tail, even when we draw right-tail observations.
- Diffuse — In 1980, you had a probability belief of
  - 1/323 ≈ 0.3% that the per-month disaster risk was 0.0%.
  - 1/323 ≈ 0.3% that the per-month disaster risk was 0.3%.
  - 1/323 ≈ 0.3% that the per-month disaster risk was 0.6%.
  - 1/323 ≈ 0.3% that the per-month disaster risk was 0.9%.
  - ...
  - 1/323 ≈ 0.3% that the per-month disaster risk was 100.0%.
- All regression results assume similar diffuse priors. With dogmatic priors, regressions are not informative.

Prior Probability Choices

- Dogmatic — the frequency (of losing all your money) is 5% (or 0.1% or 99.9%). No data can reject your prior.
- Parametric — (e.g., Veldkamp). Probability is lognormal, and we can update and compute the left-tail, even when we draw right-tail observations.
- Diffuse — In 1980, you had a probability belief of
  - $1/323 \approx 0.3\%$ that the per-month disaster risk was 0.0%.
  - $1/323 \approx 0.3\%$ that the per-month disaster risk was 0.3%.
  - $1/323 \approx 0.3\%$ that the per-month disaster risk was 0.6%.
  - $1/323 \approx 0.3\%$ that the per-month disaster risk was 0.9%.
  - ...
  - $1/323 \approx 0.3\%$ that the per-month disaster risk was 100.0%.
- All regression results assume similar diffuse priors. With dogmatic priors, regressions are not informative.

Updating Probability

- If the probability of a disaster is p, the probability of not having seen it in T draws is

$$1-(1-p)^T$$

- If the probability of a disaster is $p = 1/323$, the probability of not having seen it in T draws (months) is

$$1-(1-1/T)^T$$

- Counterintuitively, this barely / does not depend on T.
  - The probability of a 10-year flood within 10 years is 65.1%.
  - The probability of a 100-year flood within 100 years is 63.4%.
  - The probability of a 1000-year flood within 1000 is 63.2%.

Continued

- That's pretty cool, because this is

$$1 - e^{-1} \approx 63.2\%$$

- The paper contains drawings of regular and weird true population distributions. Any true population distribution can be cut up into T equal-probability segments. The probability of sampling from the lowest bin in T draws is always 63.2%.

Bayes Rule: Now We Can Update

- Given that no disaster has occurred in T draws, what is the probability that there is one that was just not sampled?

$$p(\dagger|\text{Data}) = \frac{p(\text{Data}|\dagger) \cdot p(\dagger)}{p(\text{Data})}$$

$$= \frac{100\% \cdot 1/323}{100\% \cdot 1/323 + 37\% \cdot 1/323 + 13.5\% \cdot 1/323 + \ldots} = 37\%$$

- If you grant me an equal possible frequency of disasters as an a-priori belief, with 63% probability we have already seen the worst outcome (well, a draw from the lowest 1/T bin) in the distribution. With 37% probability, we have not.
- Number of D Dark Events Worse Than Worst Observed:

None	1	2	3	4	5
63.2%	23.3%	8.6%	3.1%	1.2%	0.4%

Diffuse Prior Implications

-
  - Most likely there are no dark events if we have observed none.
  - Expected number of disasters is 0.58 (not 0, not 1)
  - X dark events is 2.71 more likely than X + 1 dark events.
  - One black swan is reasonable.
  - Five black swans are not.
  - Never 100% sure that disasters (e.g., think –100%) are not possible, even after a million draws.
  - but per-period allowable disaster probabilities would have to get smaller and smaller.
- I think an equal prior is, if anything, conservative. It is hard to imagine that your expectations were that you would see terrible outcomes half the time, just as seeing the terrible outcome once.
- An assumption of a maximum frequency can give an upper bound of the loss magnitude before even a risk-neutral investor would never purchase unprotected stock.

Long-Term Investment and Risk-Aversion

- With disaster fear, the gambler's ruin can make it (more) in the interest of a short-term risk-neutral investor to put money into stocks when a long-term investor would not.
- Not the case here, but we need to realize that the time-horizon matters, even for risk-neutral investors.

Firm Up Intuition And Perspective

- With prior, we can now use the option price to infer also maximum magnitude of expected permissible crash.
- Yes. We can assume that risk-neutrality is the least risk-aversion we believe.

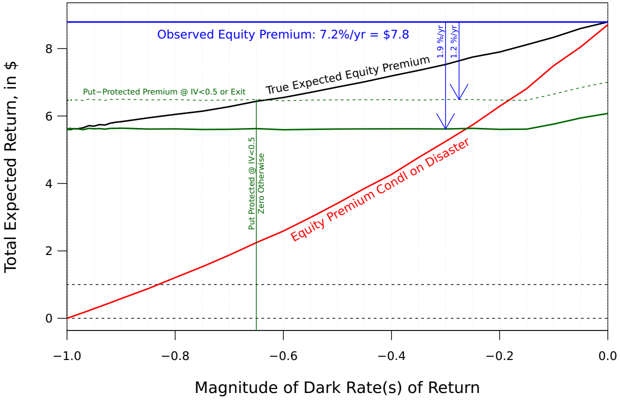

Firm Up Intuition And Perspective

- With prior, we can now use the option price to infer also maximum magnitude of expected permissible crash.
- Yes. We can assume that risk-neutrality is the least risk-aversion we believe.

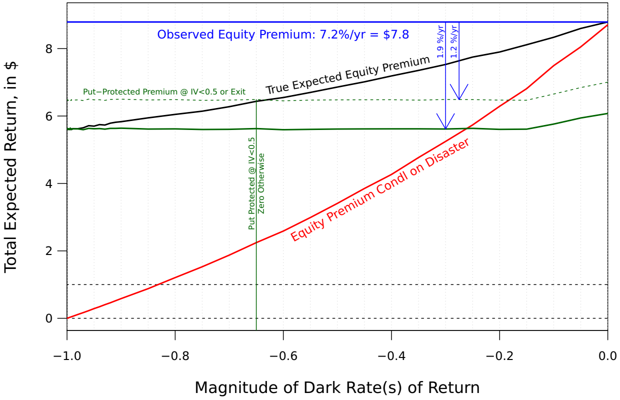

Omitted

Lots of Figures:

- About where I have option pricing data.
- How one lines up BTM options with stock (what if they trade in the morning?)
- How one works with futures that are different in expiration and possible implied dividend yield.
- How ATM option prices behaved.
- How BTM put prices behaved.
- How BTM implied vol was relative to the historical vol.

Opinion

- Disaster fears seem to have an aspect of divinity to it. It often seems impossible to test if disaster fears are rational or irrational.
  - Is this just "naming" (tautological)? Suddenly, prices move a lot. One minute, afraid; the next, not.
  - Are disasters sometimes really suddenly more likely, or are we just suddenly more afraid of them?
  - Does being more afraid of disasters make them more likely?
- Although (changes in) disaster fears "smell right," they have given some research close-to-infinite degrees of freedom to explain everything. (By definition, we did not get to see the possible disasters.)

What To Remember

- A lot of the time, disaster fears are completely unimportant. Protection is trivially cheap. Right now.
- Some of the time, it's expensive to protect oneself. Perhaps get out of the stock market then.
- About 1-2% of the 6-7% equity premium from 1983-2013 can be explained by disaster compensation. No more.
- We can model disaster frequencies, perhaps better than equilibrium prices.
- No evidence of worsening tail fears relative to ordinary 5-15% fears during the financial crisis.
- Disasters could not have been too terrible. It's hard to imagine more than a 1-in-3 chance of there having been a worse disaster what we have seen (–20% in one month), and it couldn't have been more than a loss of –50% (or put-options would have been too cheap).