FAJ
Financial Analysts Journal
Volume 72 • Number 5
©2016 CFA Institute

PERSPECTIVES

# The (Time-Varying) Importance of Disaster Risk

Ivo Welch

*How much of the historical 7% per year equity risk premium could have been risk compensation for disasters that just happened not to have occurred? The answer can be found in below-the-money put prices, which would have protected against such disasters. Using the cost of rolling over one-month index put options, I show that the maximum possible premium for crash risk could not have accounted for more than about 2% per year, thus leaving about 5% per year for reasons other than sudden disasters. I also provide a novel "conservative diffuse prior" approach for dealing with black swan risk.*

From 1983 to 2012, the geometric rate of return of the US equity premium (over short-term Treasury bills) was an impressive 7.2% a year. Can this large return be attributed to left-tail low-probability events that just happened not to have occurred? For example, if there was a "once in 200 years" catastrophic event of a 99% loss, even risk-neutral investors would have demanded 7% in non-disaster years just to break even. In my study, I investigated the potential role of such "dark" events.[1]

By definition, realized US stock returns alone cannot be easily used to assess such risk. After all, the hypothesis states that the disaster did not occur. Therefore, I looked at the pricing of deep-below-the-money index put options that would have protected against extreme left-tail events. The *maximum* possible disaster component could *not* have accounted for a premium of more than 2.1%/year because a put-protected "crash-insured" stock position (which could not have lost more than 15% in any one month) still yielded 5.1%/year. If the expected crash component had been more than 2.1%, even a risk-neutral investor would have expected to earn more by holding put-protected rather than unprotected stock—a non sequitur, however, if risk taking is compensated in financial markets. An upper limit of 2.1% is economically large but insufficient to explain the entire equity premium. To put this in perspective, even the *maximum* 2.1% disaster risk was less than the 3% sampling standard error around the mean.

*Ivo Welch is the J. Fred Weston Professor at the UCLA Anderson Graduate School of Management, Los Angeles.*

*Editor's note:* This article was reviewed and accepted by Robert Litterman, executive editor at the time the article was submitted.

The bound can be even tighter. The 2.1%/year is a "dumb" average obtained by always buying a put. Yet, the price of put insurance was time varying. At some times (e.g., during the 2008 crisis), the protection cost more than 10% a year and could have easily and fully accounted for an annual equity premium of 20%. At other times (e.g., in 2015), it cost less than 0.5% a year. Consequently, an even tighter bound on the possible influence of disasters can be obtained by considering a smarter strategy: In high-put-price months, a risk-neutral investor could have switched from put-protected stocks to Treasury bills. A simple, implementable, dynamic timing strategy—buying put-protected stocks when puts were cheap and switching the entire portfolio to Treasuries when puts were "expensive" (in absolute and volatility terms)—would also have had no exposure to left-tail disasters but would still have yielded 5.8% a year, on average. This strategy would limit the maximum influence attributable to disaster risk to 7.2% – 5.8% ≈ 1.5% a year.

It is important to recognize that rolling over deep-below-the-money put options would have provided protection, not against every kind of stock market risk but against only one particular kind of risk—that of a one-month crash. Earthquakes provide a good analogy. The California Earthquake Authority (CEA) sells insurance policies with finite time horizons (usually one year) and deductibles (usually 15%) that do not protect against every kind of harm from earthquakes. First, they do not protect against consecutive modest earthquakes, in which annual repairs repeatedly cost 15% of the home's value. Thus, they do not protect against, say, total losses of 80% over 10 years if earthquake losses arrive in many –15%/year pieces. (Such a scenario seems unlikely for a stock market price process that economics forces into a near martingale.) Second, rolling over one-year policies does not protect against higher insurance policy prices in later years, especially if a few earthquakes occur. Third, these policies do not protect against cases in which the insurance seller (the CEA) will be unable to make good on its policy promises.

In my study, I attempted to describe the implications of equivalent disaster protection for a marginal investor holding the stock market. I investigated the historical cost of rolling over one-month put option insurance with a 15% deductible, assuming that the policy sellers would have been able to make good on their promises if a major stock market collapse had occurred. Under these assumptions, the price of below-the-money index puts can decompose the equity premium into the previously mentioned maximum possible premium that could have been crash risk (about 2%) and everything else (at least about 5%). Other plausible candidates that might help explain the 7% observed equity premium include sampling variation, "ordinary" utility risk, a declining risk premium, long-run risk, and regime change.

The 1.5%–2.0% range is the product of a probability of a disaster and the magnitude of a disaster. If we are willing to adopt a "conservative diffuse prior" sampling model, we can disentangle the two. Intuitively, this article asks, "What can we learn from a long history in which we have never seen an event?" It quantifies that it is implausible to believe either that unseen events are frequent or that they are impossible. With such a model, it is then possible to assess the maximum magnitude of disasters as having been no worse than about –70%.

## Historical Index Put Pricing

For my study, I obtained data on index put options from the Chicago Mercantile Exchange (CME) for September 1983–October 2012 (323 months, or about 27 years), when CME index put options were readily available (with some missing months early on). I used Ken French's factor premium dataset to easily replicate stock returns.[2] In the 323 months for which I had good CME option-pricing data, the arithmetic monthly mean rate of return was 0.67%, the geometric return was 0.58%, and the standard deviation was 4.3%.

We want to assess the historical pricing of deep-below-the-money put option protection, but such index options did not trade frequently. In many cases, the reported end-of-day, below-the-money (BTM) put prices were from trades in the morning, whereas the end-of-day index futures prices were from afternoon trades. Calculating prices and implied volatilities from morning option prices against end-of-day futures prices as an underlying base could produce misleading results. Matching dividends and exact intraday times is important for accurate estimation of implied volatilities (used only for smoothing multiple put prices). Appendix A describes the construction of the option-pricing data in more detail.

Ideally, we would consider put-protected stock positions using puts that were far below the money—say, 30% below the money. Such options would offer good protection against –30% to –100% disasters for a very small absolute cost (annual yield drag).[3] **Figure 1** shows that such options were not actively traded. It plots the availability of low-strike-price options—specifically, the moneyness of both the fifth-lowest-strike-price option transaction and the fifth quantile's lowest-strike-price option transaction on each day. Before 1985, transactions of 15%-below-the-money put options were sporadic. But after 1985 and especially after 1987, they became more commonly available. Therefore, I focused my analysis on months with solid end-of-prior-month data on put options.

**Figure 2** shows the Black–Scholes implied volatilities of options with a strike price at the money and the *incremental* implied volatilities of put options with a strike price 15% below the money—a measure of the steepness of the volatility smile. There are both low- and high-frequency variations in these two time series, but remarkably, the steepness jumped after the 1987 crash and has remained steady ever since. The implied deep-left-tail volatility has behaved similarly to center volatility, even in 2000 and 2008. The increase in ordinary at-the-money volatility risk seems to have fully captured the increased fear of tail-risk events. Investors seem to have feared tail risk no differently than they feared ordinary volatility risk before and after the Great Recession of 2008.[4]

**Table 1** considers only options that match stringent criteria: options that were very close to 15% below the money and very close to 30 calendar days until expiration. Only about 1,200 put options had strikes of 15.5% to 14.5% below the money and between 26 and 34 calendar days until expiration. For these options, Panel A of Table 1 shows that the implied volatility was about 0.32. The inference is similar if the sample is restricted to medium-volatility periods—that is, periods in which the *preceding* three-month observed annualized log volatility was between 16% and 24%. In my data analysis, I relied on the volatility-smoothed prices of these options.

In sum, the empirical data suggest that an implied-volatility pricing of 30% (15 pps above the prevailing volatility) seems reasonably representative for the average cost of 15%-below-the-money, 30-day index puts from 1986 to 2012.

## Figure 1. Availability of BTM Index Options

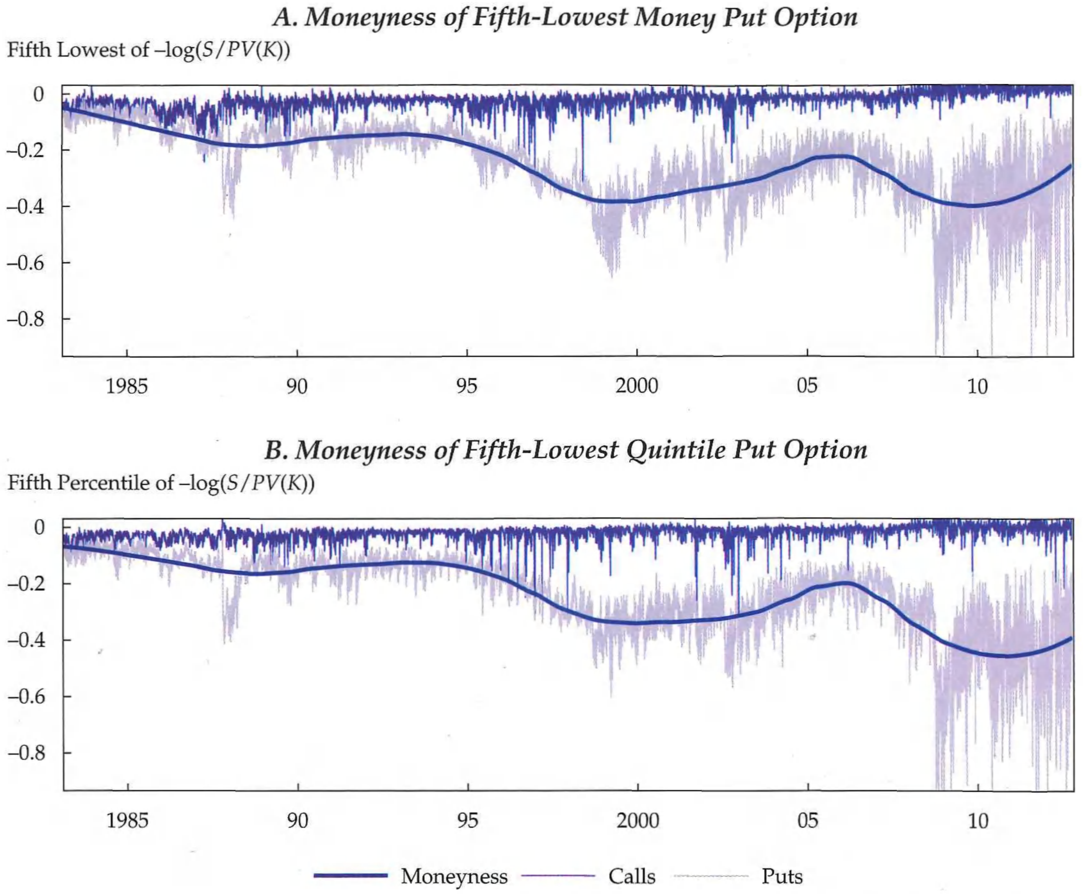


*Notes:* This figure plots the fifth-lowest money and fifth-lowest quintile put option moneyness. Put options that were about 15% below the money were irregularly available after 1983 and became generally available around 1988.

## Figure 2. Time Variation in Ordinary Volatility and Tail-Risk Volatility

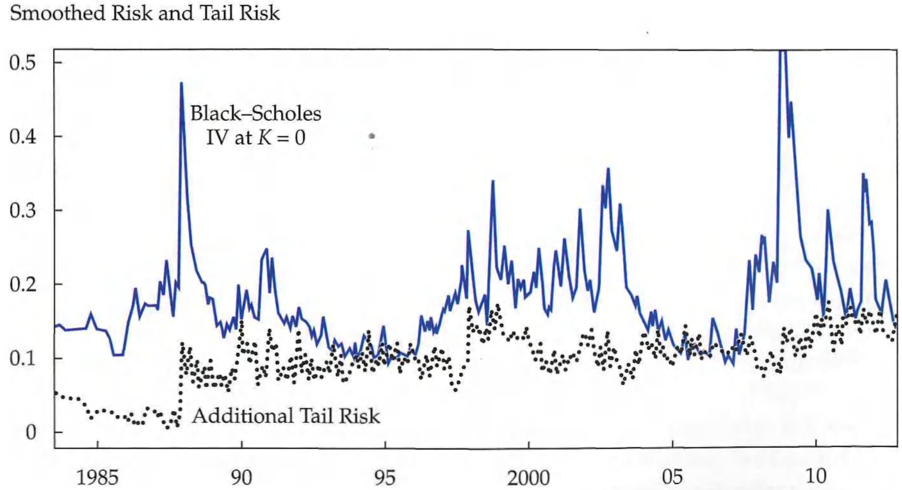


*Notes:* The solid line represents the Black–Scholes volatility calculated from options with strike prices around 0%. The dotted line represents the Black–Scholes volatility calculated from options with strike prices of –15% moneyness minus the equivalent calculated from options with strike prices of 0%. It measures the *steepness* of the left side of the (put-pricing) implied-volatility smile. There was time variation in the costs of options, both high frequency and low frequency. The steepness of the smile increased in 1987 but remained stable thereafter—including in 2000 and 2008. IV is implied volatility.

 

**Table 1. Most Applicable Monthly –15% Put Options**

| | Mean | Std. Dev. | 25th | Median | 75th |
|---|---|---|---|---|---|
| | 1,200 Options: Moneyness (–15.5%, –14.5%); Calendar Days Left (26, 34) | | | | |
| *A. 26–34 Days to expiration and –15.5% to –14.5% strike price* | | | | | |
| Annualized 3-month prevailing log volatility | 0.19 | 0.09 | 0.12 | 0.17 | 0.22 |
| Option-implied log volatility | 0.33 | 0.07 | 0.28 | 0.32 | 0.36 |
| Difference | 0.14 | 0.06 | 0.12 | 0.15 | 0.17 |
| Date | | | 1997 | 2000 | 2006 |
| | 470 Options: Moneyness (–15.5%, –14.5%); Calendar Days Left (26, 34); Historical Volatility (0.16, 0.24) | | | | |
| *B. 26–34 Days to expiration and –15.5% to –14.5% strike price, excluding outliers* | | | | | |
| Annualized 3-month prevailing log volatility | 0.19 | 0.02 | 0.17 | 0.19 | 0.21 |
| Option-implied log volatility | 0.34 | 0.05 | 0.31 | 0.33 | 0.36 |
| Difference | 0.15 | 0.05 | 0.11 | 0.15 | 0.18 |
| Date | | | 1999 | 2000 | 2007 |

## The Maximum Disaster-Based Equity Premium

We can now assess the actual performance of put-protected equity strategies during the 323-month (27-year) sample, when below-the-money CME index put options on the S&P 500 Index were readily available.

**Disaster-Protected Equity Performance.** Consider an investor who would have purchased 323 one-month CME index puts with strike prices of about 15% below the money. As mentioned in the introduction, although such puts would not have protected stock positions against (consecutive) monthly –14% rates of return, they would have protected quite well against sporadic giant disasters. There is good evidence that stock returns are not and cannot be very serially correlated because stock prices should incorporate the implications of repeated catastrophes not piecemeal but instantly.[5]

Compared with a naked stock portfolio, a put-protected stock portfolio would have been much less affected even by the most extreme possible disaster of –99% (a near-complete collapse), losing at most 15% in any one month. From 1983 to 2012, if exactly one disaster had occurred in one more month, the observed 323-month geometric mean of 0.58% would have been about

$$\left[(1+0.58\%)^{323}\times(1-15\%)\right]^{1/324}-1\approx 0.53\%.$$

Although one additional 15% loss would have hurt, it would have altered the geometric mean performance of the put-protected stock market position by "only" 5 bps per month. Considering the longer 1926–2015 sample, in which not a single –99% rate of return (thus a 15% loss) was observed, adding one would have reduced the holding rate of return on the put-protected portfolio by only about 2 bps per month. Because the probability of more than one or two terrible disasters *when not even one was in fact observed* is likely to be low (discussed later in the article), the expected reduction in the geometric mean return of the put-protected portfolio due to disasters would have been modest.

Instead, the main cost of disaster protection would have been the cost of purchasing the index puts every month. To assess the prevailing implied volatility at the start of each month, I used the implied volatilities of index put options with a strike price of about –15% over all trading days after the 23rd of the preceding month. The first month of data was April 1983 and the last month was October 2012, with some missing months (early on) in between.

**Figure 3** plots the annualized cost of monthly puts and the 12-month moving average (see also Martin 2011). The 30-day put price varied from $0 to $0.028, with a mean of $0.00125, a median of $0.0004, and a standard deviation of $0.0026.[6] There are long periods in which disaster protection was almost free (e.g., 2015) and other periods in which it was very expensive.

The total cost of rolling this monthly 15% BTM put protection over the entire sample was less than 2% a year: the key result of my study.[7] The kind of sudden stock market collapses worse than 15% in any one month could not have accounted for more than 2% of the 7% equity premium because insurance against such events could have been purchased by giving up a 2% equity premium.

**Improved Disaster-Protected Equity Performance.** Figure 3 shows that there were also periods in which the monthly put protection was

**Figure 3. Annual Average Cost of Monthly –15% Put Protection**

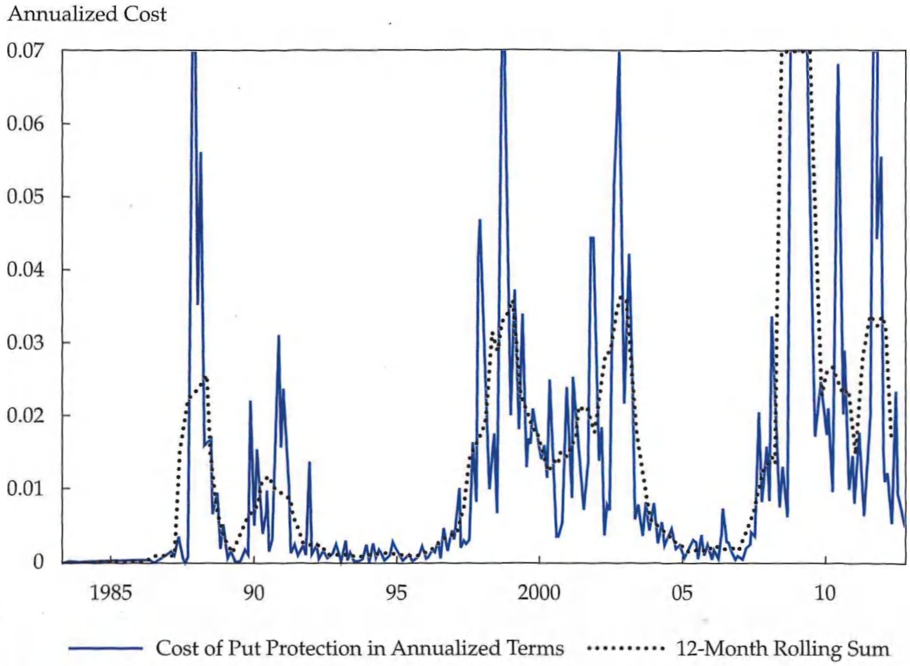


*Notes:* The solid line represents the cost of monthly –15% put protection (annualized). The dotted line represents the 12-month rolling sum. There were long stretches when disaster protection was almost free but also long stretches when it was fairly expensive.

considerably more expensive than the average monthly geometric equity premium estimate of 58 bps. For example, when the implied volatility was 50%, a one-month put with a strike price of 15% BTM would have cost about 85 bps, which raises the interesting question of whether put protection was cheap only when there was no equity premium to be had. Although we cannot measure the frequency of (unrealized) disasters in these periods, we can measure the realized *ex post* average mean excess rate of return on stocks, as shown in **Table 2**.

At times, disaster probabilities or fears were high enough that insuring against them would have been very expensive, and the strategy for "unimproved" disaster-protected equity performance would nevertheless have blindly purchased these expensive puts. The equity return in months with put costs of more than 2%/year was not only very risky but also very good.

The original question of bounding the disaster component in the equity premium now becomes the question of whether a small risk-neutral investor could have done better with a safer strategy. This investor could have decided to be disaster protected in the stock market, earning a premium of 5%–10% when put protection was cheap and exiting the stock market altogether for Treasuries when protection was too expensive (losing out on some *ex post* spectacular returns, however). Exiting the stock market altogether when put protection became too expensive would not have required sophistication.[8]

**Table 3** reports the rate of return on a strategy of put-protected stocks when the implied volatility was less than $x$ and T-bills otherwise.

All these strategies were safer than the unprotected stock investment that would have yielded 7.2%. Their returns take into account that an investor who was not in the stock market would have lost spectacular positive returns in those months when puts were expensive. The dynamic strategy that exited the market at put prices of 50% implied volatility would have yielded a geometric rate of return of 5.8% a year, suggesting either that investors earned at most 1.5% as a disaster premium (on the 7% equity premium) or that puts to insure against these disasters were too cheap even for a risk-neutral investor.[9]

**Table 2. Realized *ex post* Average Mean Excess Rate of Return by Put Cost**

| Put Protection Cost (annualized) | < 0.5% | 0.5%–1% | 1%–2% | 2%–4% | > 4% |
|---|---|---|---|---|---|
| Realized geometric equity premium | 10.5% | 5.3% | 5.1% | 18% | 13% |
| Number of months | 161 | 203 | 55 | 35 | 30 |

**Table 3. Rate of Return for Put-Protected Stock Market Strategies with Implied BTM Volatility below *x* and T-Bills Otherwise**

| | Exit Put-Protected Stock If Implied BTM Volatility Greater Than | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 0% | 30% | 35% | 40% | 45% | 50% | 55% | ∞ (never) |
| Average geometric excess return | 0.0% | 3.0% | 3.9% | 5.2% | 4.7% | 5.8% | 5.6% | 7.2% |

**Observations and Caveats.** Without additional assumptions, it is impossible to distinguish between a sudden (rational or irrational) increase in disaster fear and a sudden increase in disaster probability. Thus, it is impossible to determine what a truly risk-neutral smart marginal investor should have held.

There were months in which the put prices were so high that they effectively offered no meaningful upper bound on the importance of (compensation for) crash risk. For example, we cannot say whether equity premiums in 2009 were primarily due to compensation for disaster risk. In 2015, the bound was tight, and it seems prima facie implausible to argue—and still implausible today—that fear of disaster risk could account for a premium of more than about 1%/year.

Although I relied on actual CME market quotes and transactions, my study still had to rely on assumptions. I still used index put option price estimates, not actual transactions by an investor who followed the put protection strategy. I still assumed that investors could have purchased 30-day put options at the beginning of the month—although index options did not expire at the end of the month but, rather, at the end of the third week. I did not use the *ask* price on the option, nor did I assess the hypothetical price impact of put purchases. (I can thus report only the marginal cost of protection, not the average cost of protection if, for example, a giant fund like CalPERS had attempted to buy insurance for its whole portfolio. I did not analyze a general equilibrium model in which investors would behave differently wholesale.) I still assumed that investors could have purchased the appropriate put options at an interpolated price of similar options in the last week of the preceding month. But I also assumed that a real-world investor would not have been more strategic about buying cheaper options. (The strategies considered here did not opportunistically select the cheapest put available.) An investor could also have taken advantage of put protection at 25% below the money (when available). Judging from the (admittedly scant) evidence reported earlier, either action could have been considerably cheaper for rolled-over extreme-event protection.

Despite these caveats, there is no reason to believe that the assumptions and strategies presented here are misleading. It is inconsistent to believe both that there is a great deal of highly compensated crash risk and that put option protection against crash risk could be purchased for 0.2%/year in 2015. Either crash risk was not that important *or* below-the-money puts were significantly underpriced.

## The Probability and Magnitude of a Disaster

Does the 1%–2% disaster risk premium reflect "bad and rare" or "terrible and extremely rare" scenarios? Should we think of disasters as one-in-a-thousand or one-in-a-million probability events? Should we think of them as –100% complete losses or –50% crashes? Can we disentangle probability and magnitude? This task requires further assumptions. One common method is to impose a model of utility and/or asset pricing. In this article, I introduce another method. This method makes assumptions about priors and then estimates the frequency of disasters, in effect asking, What should we learn from historical data in which no such event was ever observed? It quantifies the intuition that we should consider events that have never occurred to be neither frequent nor impossible.

To obtain a posterior probability, we must start with a *prior* probability *before* data become available. Let us assume that this prior probability about the frequency of disasters in the population is *diffuse*. Before any data are sampled, the subjective imposed belief is that there is an equal probability that zero, one, two, or more disasters could happen. This assumption is very conservative: The probability of no disasters is no higher than that of a disaster every month. Note that this diffuse prior on disasters is not the same as assuming a constant per-period probability of a disaster (in which case, $T \times p$ disasters would be more likely to be observed). It is also not the same as assuming that the population and sampling distributions are identical.

As with all (Bayesian) analysis, a different prior would yield a different posterior. The diffuse prior starts with the assumption of a large frequency of disasters *before* any data are observed. The data will shift the prior toward a posterior with more mass on the left (few disasters). A prior that considers disasters rarer from the outset would only strengthen this inference.[10]

In this section, I show that the implication of this diffuse prior is that in the absence of even a single disaster observation in the sample, there is a 37% probability that there was/is at least one such disaster lurking in the true population distribution. Even more intuitively, it is analogous to a statement that the probability of observing an $x$-year flood within $x$ years is about 63% (at least for $x$ greater than a handful). In the stock market application, the equivalent statement is that the probability that a true "100-year financial market disaster" was not observed within the last 100 years of the sample is 37%.

**Prior Probability.** A dogmatic prior would assume parameter certainty, so there would be no learning from data. For example, if the prior is that there is a 0.0001% (or 90%) probability in each month of a –99% disaster, then no amount of data can alter this probability assessment. A dogmatic prior not only seems implausible, but it also precludes all learning from data.

Another type of prior would be to assume that the population distribution has a particular shape. For example, in Orlik and Veldkamp (2014), the prior is that the population distribution is lognormal. Data then update the parameters—here, a mean and variance. Leaning heavily on the precise distributional shape allows updating the estimate of the probability mass everywhere, including in the left tail.

Instead, we can start with a "nonparametric" diffuse prior. Let us assume that investors (in 1983) assessed the probability of zero disasters to be the same as the probability of 1 disaster, the probability of 2 disasters, the probability of 3 disasters, and so on—or even the probability of 323 disasters. The probability ($p$) of a disaster in any one month was thus

$$\begin{aligned}\text{Prior prob}\left(p=i/T\right) &= 1/T\\ &= 1/323\ \left[\forall i\in(0,\ T)\right].\end{aligned}$$

Note that this distribution is discrete and not continuous.[11] The choice of such a diffuse prior is defensible and arguably pessimistic (conservative), but it is not uncontroversial. In its defense, diffuse priors are usually implicit in *all* empirical analyses.

**Bayesian (Conditional) Updating.** Next, we need to calculate the conditional probability of observing at least one disaster when sampling, given an underlying true probability $p$ of a disaster in any one draw. For example, if the true $p$ were 0%, there would be zero possibility of drawing a disaster in any one month—and thus in any of the 323 months. The probability of seeing no disaster would be 100%. If the true probability of observing a disaster in one month were $p = 1/323 \approx 0.3\%$, there would be a $1 - (1 - 1/323)^{323} \approx 63.3\%$ probability of sampling at least one disaster in 323 months. If the true $p = 2/323 \approx 0.6\%$ in any one month, there would be a $(1 - 2/323)^{323} \approx 86.6\%$ probability of sampling at least one disaster in 323 months. And so on. Appendix A develops the conditional probabilities in more detail and provides further intuition.

Note that this updating works for *any* true population distribution, which determines where each bin starts and ends. With more data samples, the magnitude of the worst possible outcome goes down (becomes more extreme) and the size of each bin shrinks. In our case, think of dividing the assumed true probability distribution into 323 bins, each with an equal probability mass of $1/323 \approx 0.3\%$. Assume that the true population distribution is such that the leftmost bin, containing probability mass 1/323, ranges from –100% to –99% and that the next leftmost bin ranges from –97% to –99%. Then sample from the population distribution 323 times. The probability of drawing at least once from the leftmost bin (i.e., given this true distribution, a rate of return worse than –99%) in 323 tries is about 63.3%. The probability of drawing at least once from the two leftmost bins (i.e., observing at least one rate of return worse than –97%) is 86.6%.

**Posterior Probability.** Using the diffuse prior and the conditional probability of observing a disaster, Bayes' rule gives the posterior probability of dark events in the population. For example, the posterior probability that there are truly zero disasters in the population ($p = 0$), given that none were observed in 323 sample draws, is

$$\text{Posterior prob}\left(p=0\right)=\frac{100\%\times 1/323}{100\%\times 1/323+63.3\%\times 1/323+86.6\%\times 1/323+\cdots}\approx 63.2\%.$$

Using some algebra, we can show that the marginal and cumulative probability density functions generalize to

$$f\left(i\right)\approx\frac{e-1}{e^{i+1}},\ \ F\left(i\right)\approx 1-e^{-(1+i)}\quad \forall i\in\left(0,1,\cdots\right),\ (1)$$

where $i$ is the number of disasters $\left[i=0\Rightarrow f\left(0\right)=F\left(0\right)=62.3\%\right]$. Given that no disaster occurred (i.e., in our example, no rate of return of –99% or worse was observed), the posterior probability, as shown in **Table 4**, is 63.3% that $p = 0$ (no probability of a rate of return below –99%), 23.3%

**Table 4. Probability of Number of Dark Events Drawn in True Population Being Worse Than the Observed Worst Draw**

| | 0 | 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|---|
| Probability | 63.3% | 23.3% | 8.6% | 3.1% | 1.2% | 0.4% |

 

that $p = 0.3\%$ (a 1/323 probability mass below –99%), 8.6% that $p = 0.6\%$ (a 2/323 probability mass below –99%), and so on.

It is counterintuitive but correct that the posterior probability function does not depend on the number of samples ($T$) to a first-order approximation.

In sum, sampling fairly from a stable underlying distribution, a diffuse prior over unobserved dark events yields a posterior probability with some appealing features:

1. The sampled distribution is the maximum likelihood estimate of the population distribution.
2. The probability that we have sampled any one population bin (specifically, the leftmost bin) in a sample is about 63.2%, which is about 2.7 times as high as the 23.3% probability that a specific dark event has not yet been observed.
3. Although it is quite plausible that there are some unobserved dark events, it is outright implausible to assume that five or more dark events (and specifically the leftmost five) have never been sampled.
4. There is about a 2-in-3 chance that there was no lurking, unobserved dark rate of return (i.e., that we have already sampled from a particular probability mass, such as the leftmost one) and a 1-in-3 chance that there was.
5. The probability of one or more dark events is $1 - f(0) \approx 36.8\%$.
6. The expected number of unobserved dark events is not zero but, rather, about $\sum f(i) \times i \approx 0.58$.

## The Maximum Magnitude of Unsampled Disasters

We can now use the maximum equity premium component (1%–2%) and the estimated frequency of unobserved disasters (about 37%) to infer a reasonable maximum magnitude of a disaster in the population: Such unobserved disasters could not have been expected to be worse than –70%. If disasters had been any worse, even a risk-neutral investor would have been better off not holding unprotected stock. The frequency of such disasters would simply be too high. Again, a less diffuse and less pessimistic prior (i.e., that disasters are less common) would yield even more-modest estimates of the worst.

**Investment Horizon and Disaster Risk.** Our concern is not that there were a few unsampled dark realizations (events) in the center of the population distribution. Rather, our concern is that there were a few high-magnitude left-tail realizations that happened not to have been observed (yet). It is easiest to assume that our concern is with dark returns of one specific magnitude, $r_D$. (If the magnitudes of dark returns are heterogeneous, most of the discussion still applies to the geometric means of these dark returns.) Most bounds emerge if we consider the worst case: an unrecoverable rate of return of $r_D = (-1)$.[12]

The probability of a disaster has a counterintuitive effect on the behavior of investors: When there is a disaster probability, even risk-neutral investors cannot ignore their investment horizon. That is, an investor who plans to be in the market for 1 year would have a different trade-off than an investor who plans to be in the market for 10 years, even if the latter investor can exit after 1 year. Perhaps more surprisingly, in the presence of disasters, both "gambler's ruin" and compounding can work together so that long-term investors prefer bills and short-term investors prefer stock.

For example, if the prior is that a disaster has a dogmatic positive probability (100%)—say, a 1/323 probability of $r_D = -1$ and a 322/323 probability of $r_G = 0.0058$ for one month—then a single-month naked-stock investment would have an expected rate of return of $100\% \times \left[(322/323) \times 0.0058 + 1/323 \times (-1)\right] \approx 0.27\%$ over that month. With a positive expected rate of return, a sufficiently risk-tolerant investor would prefer to be in the stock market. As $T$ increases, however, the probability of falling into ruin quickly increases from 1/323 to 1. Thus, with a dogmatic prior, a short-run risk-tolerant investor may prefer to be in the stock market, whereas a long-run risk-tolerant investor may not.

The issue of investment horizon dependence also applies to diffuse priors. The trade-offs for a long-term investor are different from those of a short-term investor. Consequently, we can answer only the question of what the maximum magnitude of a disaster could have been to keep a risk-neutral investor in the market given a specific investment horizon. Thus, we can consider an investment horizon, $Tf$, equal to the sampling horizon, $Th$ (i.e., 323 months). The posterior expected rate of return is

$$E(r) = \sum_{i=0}^{Th} f\left[i(Th)\right] \times \left\{p(i,Tf) \times r_D(Tf) + \left[1 - p(i,Tf)\right] \times r_G(Tf)\right\},$$

where $p(i, Tf)$ is the probability of drawing $i$ disasters within $Tf$ investment periods. For a one-month investment, $p(i, Tf) = 1/Th$. For an infinitely lived investment, $p(i, Tf) = 1$ (i.e., the $D$ event would occur with certainty). Note that both $r_D$ and $r_G$ also depend on $Tf$, unless they are –1 or 0. Consider a case in which we have 323 months of historical data, the no-disaster return is $r_G(Tf) = \left[(1 + r_G)^{Tf} - 1\right] \approx 1.0058^{Tf} - 1$, and the disaster is the worst case, $r_B(Tf) = -100\%$.

Given these probability assessments of disasters, an investment in the (unprotected) stock market would have a true expected rate of return of

$$E(r) \approx 0.632 \times 0.58\% + 0.232 \times \left[\left(\frac{322}{323}\right) \times 0.58\% + \left(\frac{1}{323}\right) \times (-1)\right] + 0.086 \times \left[\left(\frac{321}{323}\right) \times 0.58\% + \left(\frac{2}{323}\right) \times (-1)\right] + \ldots \approx 0.4\%.$$

This average rate of return would have been high enough to keep risk-neutral investors in the (unprotected) stock market, even given the possibility of a complete loss.

**The 1982–2012 Sample.** We can now assess the expected rate of return as a function of the horizon for an investor who planned to remain in the market from 1982 to 2012. **Figure 4** considers the historically observed 323 months (with CME below-the-money put option pricing data) and what investment strategies a risk-neutral investor with such an investment horizon might have preferred. As noted, both the historical sample period (*Th*) and the investment horizon (*Tf*) are set to 323 months.

To arrive at Figure 4's calculations, we resample the observed no-disaster rates of return (together with historical beginning-of-month put pricing) and "inject" a disaster of a given magnitude ($x$-axis) according to the diffuse prior and specific posterior in random months. That is, we inject "no disaster" with 0.632 probability ($p = 0$), one disaster with 0.232 probability (thus drawing a disaster of magnitude $x$ in each of the 323 months, with probability $p = 1/323$), two disasters with 0.086 probability ($p = 2/323$), and so on. The $y$-axis is the expected net equity return for an investment of 323 months.

The top blue line depicts a resampled naked-equity premium without disasters. It would have accumulated to about \$8.80 with random resampling.[13] The true expected rate of return for the

**Figure 4. Dark-Return Magnitude, Safer Strategies, and Equity Premiums**

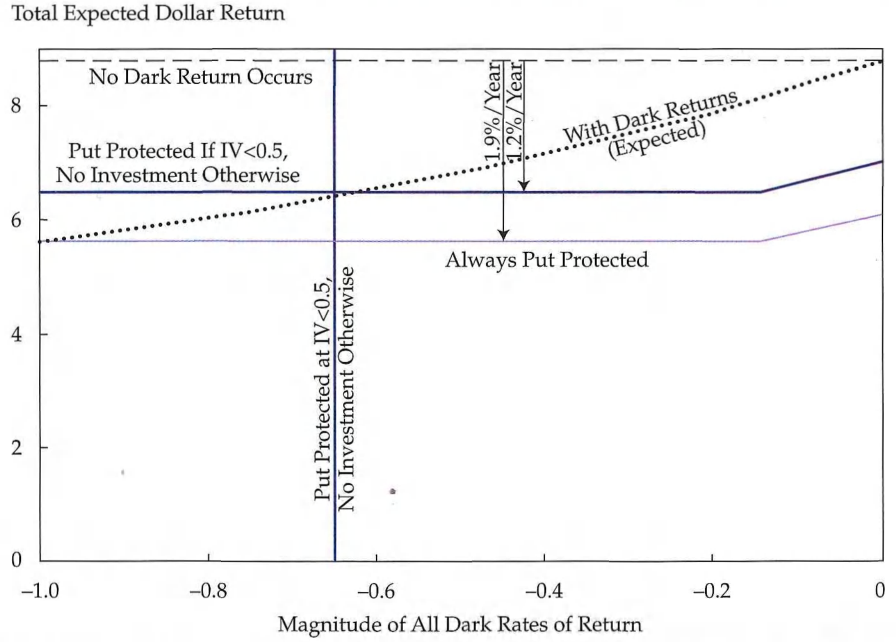


*Notes:* This figure illustrates the origin of the maximum equity premium that can be attributed to disasters (dark returns). Resampled, if there are no dark returns, unprotected equity in the absence of disaster would have yielded 7.2%/year. By the thought experiment, we have been living in just such a world, which happened not to have experienced even one dark return—represented by the dashed line. Now, without loss of generality, assume that all dark returns are the same. Further assume that they occur with a frequency equal to the diffuse posterior distribution—for example, 23.3% for one dark return, 8.6% for two dark returns, and so on. The dotted line represents the true expected rate of return. Finally, the put-protected equity strategies have limited losses—even if the stock market were to lose, for example, all its value. An equity strategy that is always put protected would yield (at least) just under \$6. An equity strategy that invests in put-protected equities if the implied volatility is less than 50% and exits all financial markets otherwise would yield (at least) just over \$6. Given probabilities, if the dark-event return were any worse than –65%, the true expected rate of return on stocks would be less than the true expected rate of return on put-protected stocks. No risk-averse or risk-neutral investor should purchase unprotected stock any longer—a non sequitur.

 

naked-stock position is determined by the disaster probability assessment (the diffuse prior) and the forward-looking investment horizon (323 months), which creates the dotted line. For example, if the unknown disaster had been a one-month realization of –80%, the true expected payoff would have been about $6, or $2.80 less than the $8.80 resampled return.

The point of intersection of the put-protected stock position and the true expected equity premium fixes the disaster's maximum magnitude. If an investor had protected the stock with dynamic puts, the expected maximum magnitude could not have been worse than about –65%. If it *had* been worse, even a risk-neutral 323-month investor would have been better off protecting the stock with puts instead of holding naked stock.[14]

There are two more noteworthy observations in Figure 4. First, it illustrates nicely why the curvature of the true equity premium (determined by the diffuse prior) does not affect the calculation of the equity premium's disaster component, which is the distance between the put-protected and the naked-stock positions. The static put-protected portfolio would have ended up with about $5.60 rather than $8.80, for a geometric drag of about 1.4% a year. The dynamic put-protected portfolio—exiting both the stock and the put if the implied-volatility pricing of the put went above 50%/year—would have ended up with about $6.50, for a geometric drag of about 0.8% a year. Regardless of where on the $x$-axis the average rate of return on the true naked-equity premium intersects the put-protected equity premium, the distance between the lines is the same.

Second, we would have calculated an expected rate of return of zero if one disaster of magnitude –85% had occurred. We would probably be confused even if we had seen a disaster much less worse than –85% (not shown in the figure): We would then fail to reject the null hypothesis of risk neutrality, even if the disaster was "only" –50%. Disaster risk premium models operate in a very narrow band, with an ability to lower the equity premium from the historically observed, seemingly large geometric 6%–7% a year to the more model-pleasing 3%–4%, on the one hand, or to a nonrejectable zero equity premium, on the other hand. The reason is that if no disaster is observed, the return seems too high, but if just one disaster is observed, it seems too low and can no longer be distinguished from zero at conventional statistical significance levels. If it were ever observed, the journals would probably have to publish papers in which authors would have to design clever tests to reject the hypothesis that the world is risk neutral.

## Conclusion

In this article, I have documented the following:

1. Over the sample period, far-below-the-money put options were cheap enough that dark returns cannot explain more than a maximum 1%–2% component of the (6%–7% geometric) equity premium. This magnitude would seem to be nontrivial, but it is not even as large as ordinary sampling variation (3%).
   For an investor, this finding means that a drag of about 1%–2%/year was the cost of escaping the risk of extreme crashes. At most, a third of the superior stock market performance in the modern era can be attributed to extreme dark events—very large black swans. The rest must have been earned for other reasons.
2. The cost of put protection was time varying. At times (e.g., in 2008), crash fear could have accounted for more than the average 6%/year realized equity premium itself.
   At other times (e.g., in 2015), crash fear could not explain more than a 0.5% equity premium per year, because investors could buy such "cheap" crash protection.
3. The volatility smile did not steepen during the financial crisis of 2008. Extreme left-tail put protection did not increase in price more than at-the-money put protection.
   For an investor fearful of black swans, even 2008 was not special. All risk premiums increased, but black swan risk did not increase even more. There was nothing special about extreme dark-event risk in 2008.
4. Although zero dark-event disasters are the scenario of maximum likelihood, given the return history of the modern era, there is still a 37% probability that more black swans exist in the true population even under diffuse conservative priors.
   For an investor, it is not irrational to believe in the possibility of an unprecedentedly terrible stock market crash, despite many decades in which no such crash has occurred.
5. A diffuse prior bounds the maximum magnitude of a disaster to about –60% to –80%. Of course, this is an upper bound; the crash magnitude may well be much more modest.
   For an investor, it is irrational to believe that a black swan will wipe out the entire stock market in one fell swoop. The market pricing of options rejects such a belief.

Why have black swans received so much attention, in both the popular press and the academic press, compared with, say, ordinary sampling uncertainty? One answer may be time variation. Disaster fears and probabilities are evidently more important in some years than in others. Another, more cynical answer is that disasters have often been conjured rather than measured, which gives great freedom to speculate about all sorts of financial phenomena: Investors could fear real or imaginary disasters one day but not the next. Without an empirical measure, given the extremely low incidence of black swans, such claims have been hard to refute. Disasters have become a deus ex machina—a theology of the divine to explain mysteries that cannot be understood by other means.

In this article, I have suggested that price data on far-below-the-money puts can be used to measure time variation in crash risk both effectively and quantitatively. Daily data on the cost of disaster insurance have been available since the mid-1980s. The data used by me are available as supplemental material at www.cfapubs.org/doi/suppl/10.2469/faj.v72.n5.3.

It is, of course, true that these data are not enough to disentangle time-varying fear of disasters from time-varying disaster probabilities, but the data can help us understand whether disaster models can explain at least some of the time variation in the prices of far-below-the-money put insurance. Although disaster models that cannot explain the time variation are possible, they do not seem plausible to me. They would have to rely on a strange negative correlation between fear and probability.

I leave the reader with two crucial questions: Is large crash risk really not that important, or are put options too cheap? If not crash risk, then what can explain the historically high equity premium?[15]

*An early version of this article circulated as "Some Quantitative Limits for Disaster Risk and Equity Premium Estimates" (April 2013 and June 2014). I am grateful to Luis Garcia-Feijóo, who invested a great deal of time, effort, and good judgment to improve this article, and to Barbara Petitt, Ian Cooper, and one anonymous referee. They all made this version much better than my original submission.*

*Editor's note:* This article was reviewed via our double-blind peer-review process. When the article was accepted for publication, the author thanked the reviewers in his acknowledgments, and the reviewers were asked whether they agreed to be identified in the author's acknowledgments. Ian Cooper was one of the reviewers for this article.

CE Qualified Activity CFA Institute 0.5 CE credit

# Appendix A. Option-Pricing Data

This appendix describes the construction of the option-pricing data. The original source data consisted of 1 million put option transactions from inception in 1983 to December 2012. The Chicago Mercantile Exchange (CME) data were difficult to work with, partly because the CME sometimes included and sometimes did not include quote data. The CME data also contained a good number of obviously incorrect transactions. I excluded options closer than 3 days to expiration (when closing-out transactions could create anomalous prices) and options with more than 200 days to expiration (rarely traded). The median number of days to expiration in the remaining raw data was 35 days, and the mean was 45 days. I also excluded all options that traded before 8:00 a.m. or after 4:00 p.m.

The times and dividends were carefully matched to translate index and option prices into reasonable implied Black–Scholes volatilities.

## The Original Intraday Data

The original sample, as shown in **Table A1**, of CME S&P 500 Index[16] intraday data consisted of about 42 million records, both quotes and transactions, from inception in 1983 to December 2012. As already noted, the quote data proved unreliable because the CME sometimes did and sometimes did not include them. After removing these and other obviously incorrect records, there were about 3 million option transactions (about 400 option trades per trading day) and 28 million futures transactions. The final sample contained 822,260 call options and 984,987 put options.

As previously noted, I excluded all option transactions closer than 3 days to expiration (when closing-out transactions could create anomalous prices) and options with more than 200 calendar days to expiration (they were rarely traded).

**Table A1. Intraday Transaction Data**

| | Options | | Futures | |
|---|---|---|---|---|
| | Electronic | Floor | Electronic | Floor |
| Started with | 1,710,150 | 4,883,571 | 11,264,460 | 23,835,349 |
| Removed canceled trades | 1,710,085 | 4,732,868 | 11,264,093 | 23,635,275 |
| Removed cabinet trades | 1,710,085 | 4,719,214 | 11,264,093 | 23,635,275 |
| Removed zero prices | 1,710,085 | 4,719,214 | 11,264,093 | 23,629,597 |
| Removed quotes, % of data | 99% | 38% | 27% | 18% |
| Actual trades left | 4,388 | 2,908,457 | 8,265,482 | 19,459,330 |

Interpolation was required to determine the value of the market index that matched each option. (This was also why it was impossible to use the end-of-day data: The end-of-day stock market price may have moved between the option trade and the end-of-day trade—and more so for more sparsely traded below-the-money options.) Thus, I matched each option to a prevailing-at-this-second index price. Unfortunately, even the index futures did not trade every second. Moreover, index futures and index options often had different expiration dates. Therefore, for each option transaction, I interpolated a prevailing index futures price from the surrounding nearest-in-transaction-time and nearest-in-expiration index futures. If more than five minutes had elapsed since the last transaction or more than five minutes were to elapse before the next transaction, I discarded the option transaction. In the final sample, the median time from the option transaction to the last futures trade was 5 seconds and the median time to the next trade was 31 seconds. I limited the extrapolation in expiration timing to –30% and +130%.[17] The median option expired on the same day as the near S&P futures trade.

For a risk-free rate that was applicable to the price of the option, I interpolated a prevailing risk-free rate from the three-month T-bill rate and the one-year Treasury rate. This was done by finding the matching convex combination using the three-month T-bill (DTB3) and the one-year T-bill (DTB1YR), both obtained from the Federal Reserve Bank of St. Louis FRED database—again, limited to a range of –30% to 130% off the convex combination.

The S&P 500 Index does not include dividends, which creates a complication with respect to the dividends paid by the index. Although market-based methods show greater volatility than the historically prevailing dividend yields that are often used, it makes little difference which is used. I inferred a dividend yield from the spot-versus-futures price differential. The put itself pays off on the basis of the stock price net of dividends, which means that I needed to translate between index futures and index spot values. I started with an estimate of the prevailing index dividend yield net of the risk-free rate. One common choice is to use the historical dividend yield in Shiller (2005), which is readily available online (e.g., www.multpl.com or www.econ.yale.edu/~shiller/data.htm). It is also possible to use a market-implied estimate. I extracted a prevailing implied dividend yield net of the risk-free rate from the difference between the widely available S&P 500 NYSE end-of-day index value (4:00 p.m. EST) and all available index futures that traded between 2:55 p.m. and 3:05 p.m. PST—all related by the arbitrage condition $S_0 = I_0 \times \exp(\delta - r_f)$. On occasion, the futures minus spot difference suggested a negative dividend yield. Because investors cannot be forced to pay in such cases, I Winsorized the net dividend yield at zero.

**Figure A1** plots the left side of the well-known volatility smirk. The average implied volatility was about 3% higher than the prevailing historical volatility (not shown). Put options that were more below the money traded at higher implied volatilities. Figure A1 suggests that an average implied-volatility estimate of 32.5% (15% above the prevailing volatility) would be on the high side. However, options that were much closer to or further from –15% moneyness influenced the implied-volatility estimates. Thus, **Figure A2** focuses only on options that were closer to the –15% target. It plots the prevailing volatility on the $x$-axis and the average implied volatility of –13% to –17% moneyness options on the $y$-axis. There is a regression effect: The implied volatilities suggest that actual volatility was expected to mean-revert. Even when the annualized historical volatility was as low as 10%, put options that were 15% BTM still traded at about 20% implied volatility (i.e., 10% above prevailing). When historical volatility was high (e.g., 40%), this difference shrank to about 5%.

## The Conditional Probability of a Disaster

**Figure A3** illustrates a discretized form of the population distribution, with $T$ equally likely probability areas at different locations (specific equity premium magnitudes). To help with the intuition, consider the case in which the true underlying density function is normal. Divide it into 10 segments, each containing 10% probability density.[18]

Now, sample from this distribution $T = 10$ times. The probability of sampling (one or more times) from the leftmost bin—which turns out to be realizations of –1.3 or lower in the case of a normal distribution—is

$$1-\binom{T}{0}\times\left(\frac{1}{T}\right)^0\times\left(1-\frac{1}{T}\right)^T \approx 1-\left(1-\frac{K}{T}\right)^T$$

$$\approx 1-\left(1-\frac{1}{10}\right)^{10} \approx 65\%.$$

The probability of sampling (one or more times) from the two leftmost bins—which turn out to be realizations of –0.9 or lower in the case of a normal distribution—is

$$1-\binom{T}{0}\times\left(\frac{2}{T}\right)^0\times\left(1-\frac{2}{T}\right)^T \approx 1-\left(1-\frac{2}{T}\right)^T$$

$$\approx 1-\left(1-\frac{2}{10}\right)^{10} \approx 89\%.$$

## Figure A1. Implied-Volatility Smirk of Index Put Options

*A. Implied Volatility as a Function of Moneyness*

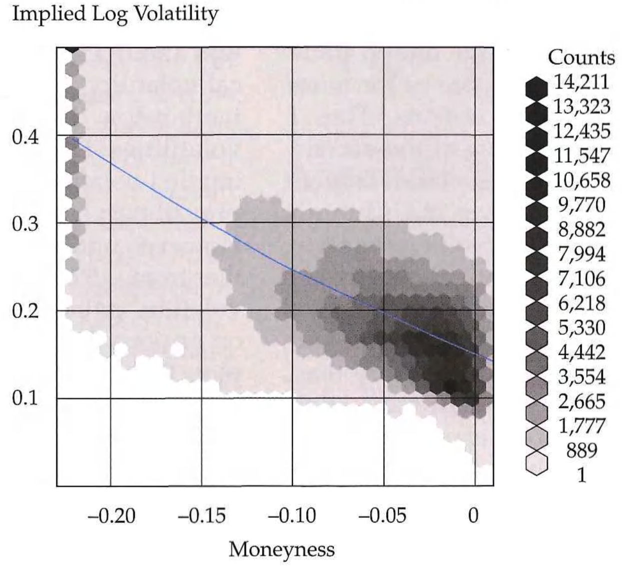


*B. Implied Volatility Net of Prevailing Historical Volatility as a Function of Moneyness*

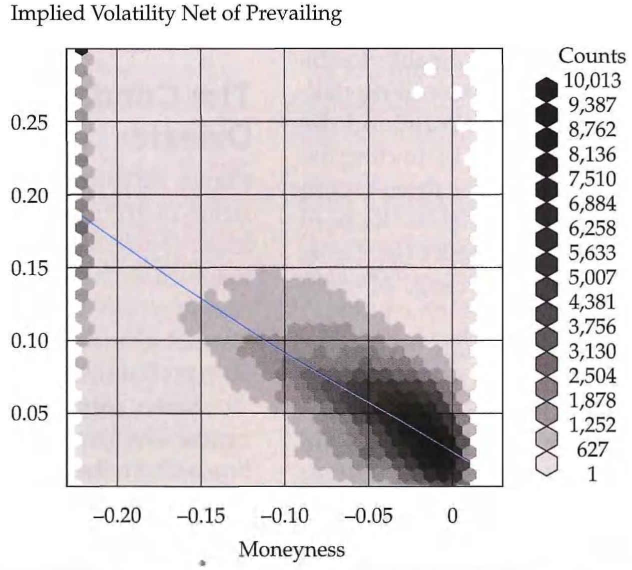


*Notes:* The $x$-axis is the moneyness; the $y$-axis is the implied Black–Scholes volatility. Unlike in the previous figures, each option (and not each transaction day) represents one observation. Panel B subtracts the prevailing three-month historical volatility from the Black–Scholes implied volatility. This figure depicts the left side of the volatility smirk; 15% BTM puts were priced at about 30% implied volatility.

Naturally, the analogous probability of drawing an impossible observation (i.e., no bin) is 100%.

The normal distribution is just for illustration. **Figure A4** shows an alternative true population distribution in which there are two very negative possible realizations, each with a probability of 5%, and nine other bins, each containing a probability mass of 10%.

The same probability calculations apply, except that the leftmost bin now contains all realizations below −4 (as drawn).[19] The probability of sampling the lowest 10% probability value at least once in 10 draws is 65%.

A remarkable property is that these probabilities are roughly invariant to the number of samples drawn because $(1-K/T)T \approx e^{-K}$ for $T/K >> 5$. Given whatever (lowest) realizations have been seen in the sample thus far, the probability that $K = 1$ more probability mass (to the left of this lowest realization) that was not drawn is the same regardless of whether 10 or 10,000 realizations have been observed, as shown in **Figure A5**. That is, if we had sampled 10 draws and divided the true distribution into 10 possible (equally likely) population areas, the probability of

 

**Figure A2. Implied Volatilities of 15% BTM Index Put Options vs. Historical Volatility**

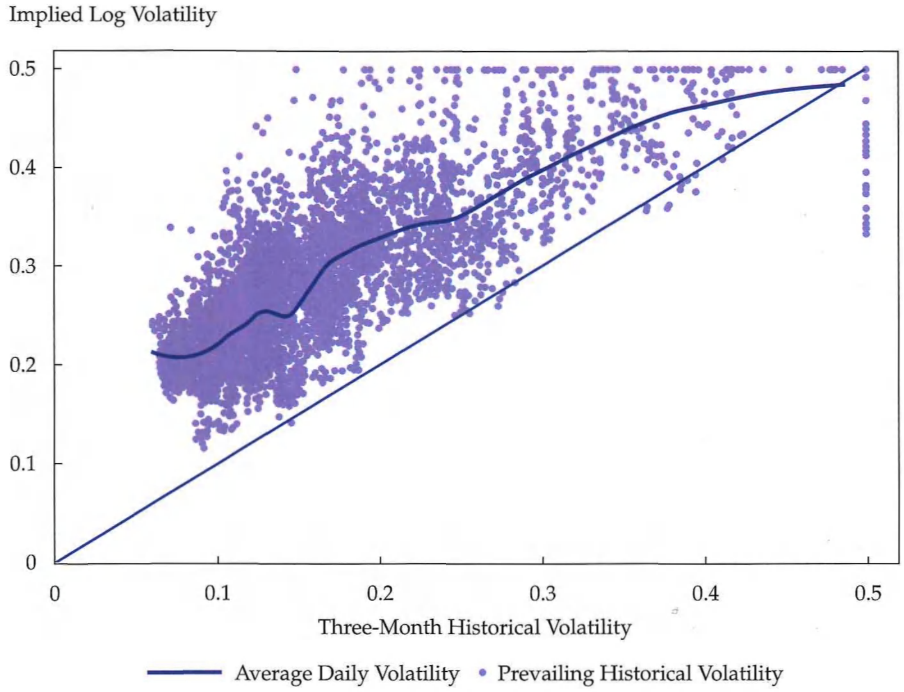


*Notes:* This figure depicts the average daily volatility for all index puts with moneyness between –13% and –17%, plotted against the prevailing historical volatility. Implied volatilities indicate volatility mean reversion.

**Figure A3. Equal Probability Slices of a Gaussian Normal Distribution**

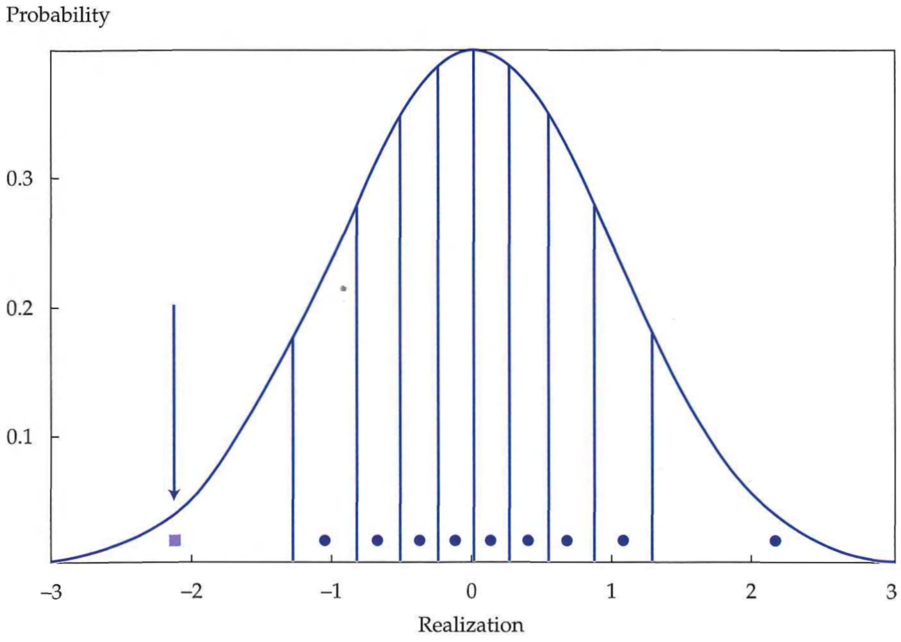

**Figure A4. Equal Probability Slices of an Arbitrary Distribution**

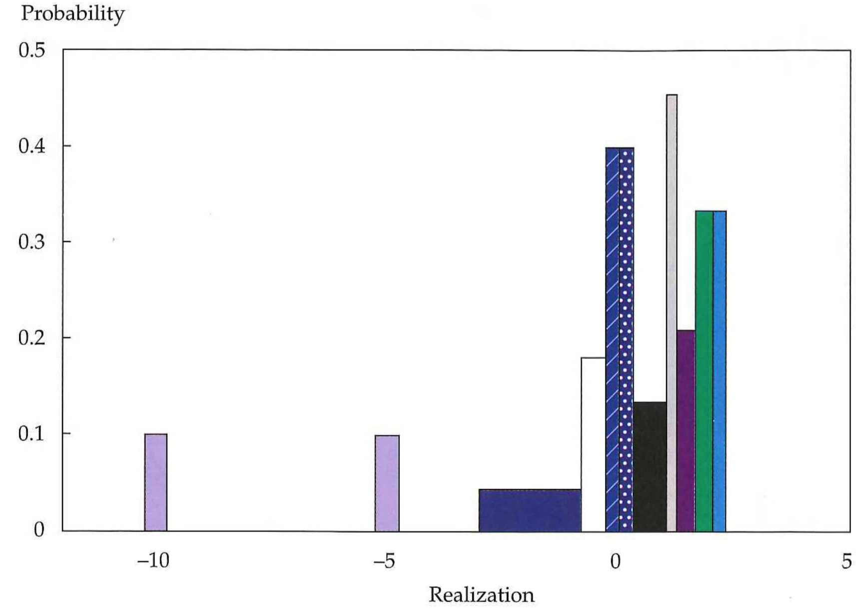


**Figure A5. More Slices of Gaussian Normal Distribution**

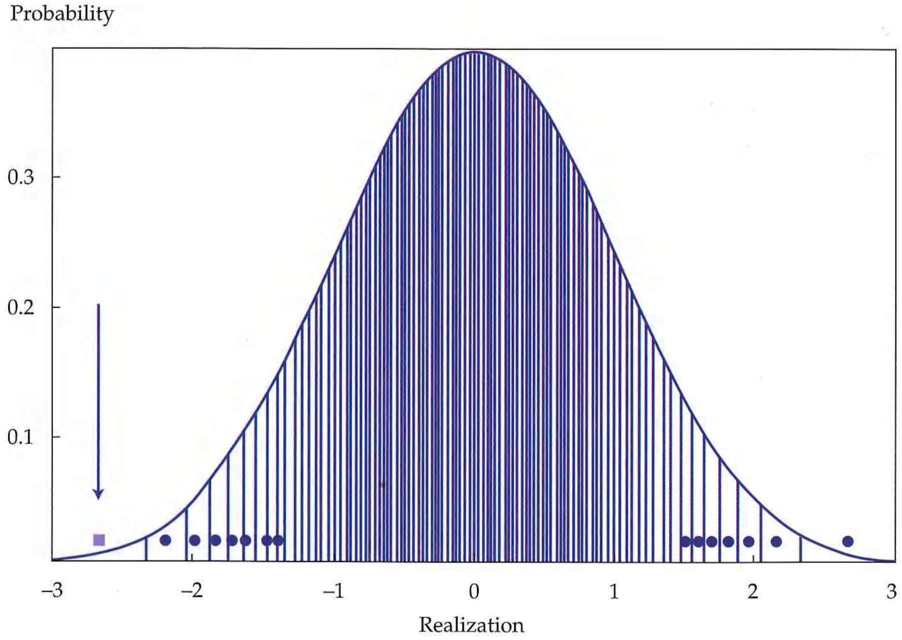


 

not having seen any particular interval (specifically, the worst) would be 36.8%. If we had sampled 100 draws (e.g., from a normal distribution in which we assigned the left bin all numbers less than –2.36, or 1% of the mass), the probability of not having seen such a number would be 36.8%.

The probability of not sampling the leftmost bin is

$$\left(1-\frac{1}{T}\right)^T = \binom{10}{0} \times \left(\frac{1}{10}\right)^0 \times \left(1-\frac{1}{10}\right)^{10}$$

$$\approx \binom{100}{0} \times \left(\frac{1}{100}\right)^0 \times \left(1-\frac{1}{100}\right)^{100} \approx e^{-1} \approx 36.8\%.$$

It is somewhat counterintuitive that this probability does not depend on $T$ (i.e., on whether 10 samples or 100 samples have been drawn), but it is a direct consequence of fair sampling. In our application, in contrast to the usual binomial formula application, $p \equiv K/T$ decreases with the number of samples drawn because the number of sample draws, $T$, also informs us about $p$. Even as we obtain more and more sample draws, we are interested not in the probability of not having seen the lowest $k$% quantile realizations but, rather, in the probability of not having seen the lowest $K$ realizations. As $T$ increases, the expected minimum draw value decreases and the probability mass within each segment decreases.

## Notes

1. Mehra and Prescott (1985) raised the puzzle; Rietz (1988), Taleb (2004), and Barro (2006) formulated the "left-tail" answer. There is also an active macrofinance economics literature about disasters (see, e.g., Gourio 2012; Wachter 2013; Gao and Song 2013; Kelly and Jiang 2014; Kozeniauskas, Orlik, and Veldkamp 2014; Orlik and Veldkamp 2014; Chen, Dou, and Kogan 2015). However, none of them used put option prices to bound disaster risk. Seo and Wachter (2015) modeled time scales of disaster risk from implied volatility curves. Coval, Jurek, and Stafford (2009) considered CDX bonds as economic catastrophe bonds, which they viewed as overpriced investments in the left tail of the market. Rosenberg and Engle (2002) found time variation in an empirically estimated kernel for 1991–95 option-pricing data.
2. See http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.
3. Although such puts tended to have high Black–Scholes volatility, they still tended to have very low prices, which is the primary aspect in this case.
4. The implied-volatility data, albeit without support, are available as supplemental material at www.cfapubs.org/doi/suppl/10.2469/faj.v72.n5.3 and my website (http://ivo-welch.info).
5. This fact does not apply to stock return volatility! Nowotny (2011) examined the Great Depression, in which volatility begat more volatility and/or possibly lower future expected rates of return. The protection cost in our example is overstated. The put would also have protected and paid off on rare occasions in the sample itself. Even under ordinary near-diffusion risk, in the absence of catastrophes, it would sometimes pay off. The put would not, however, protect against a new regime in which the expected rate of return would be very low for years to come.
6. The implied volatilities ranged from about 13% to about 77%, with a median of 26.5%, a mean of 28%, and a standard deviation of 8%.
7. Orlik and Veldkamp (2014) used nontail realizations to model time-varying left-tail disaster probabilities, leaning strongly on their distributional assumption. One could test whether their variation in implied disaster probabilities coincides with the variation in input prices.
8. One could also consider strategies that protect the stock only if puts are cheap and that remain in the stock otherwise. Such strategies are still safer than the always-in-stock strategy. For example, an investment strategy that protected against disasters only when the put's implied-volatility price was less than 20% and that remained unprotected in the market not only was safer but would also have yielded average rates of return of about 6%. Unfortunately, the key assumption that disasters were i.i.d. (independently and identically distributed) draws would no longer be reasonable: Perhaps disaster probabilities were higher when the implied volatilities were higher, and thus a dynamic put exit strategy cannot deliver a convincing bound. Until we start observing some disaster events, we cannot assess whether dark-event disasters are more frequent when put options are priced higher than 30% volatility. Thus, we cannot determine whether a 1% or 2% drag is more plausible.
9. My CME dataset ended in October 2012 and thus did not include 2015 data. In mid-2015, the cost of protection remained low:

   On 16 April 2015, the S&P 500 stood at 2,105. A put option with a strike of 1,800 expiring on 8 May 2015 (three weeks) had an ask quote of $0.35 (bid of $0.25) for an implied volatility under 20%. This put cost $0.00017 per $1 of protection. With 17 three-week periods a year, if the price of BTM volatility protection remained stable, the prevailing put price implied a cost drag of about $17 \times 0.35 / 2{,}105 \approx 0.2\%$ a year.

   On 20 April 2015, the S&P 500 stood at 2,086.20. A put option with a strike of 1,600 expiring on 15 July 2015 (85 days) had a last transaction of $0.70, for an implied volatility of 33%. Rolling over 85-day contracts 4.3 times covers the year. Thus, the cost of protection at a strike of 67% of the stock was $3 a year, or about 0.15%.

   On 24 April 2015, the S&P 500 stood at 2,118 and a June 2016 option with a strike price of 35% below the money (1,375) cost $22.80, or about 1% of the index. Such a put would even protect against repeated gradual stock market declines. In 2015, disaster protection was *very* cheap, though not in a Black–Scholes price sense. Indeed, these below-the-money index put options are still too expensive in the Black–Scholes sense; they are cheap in the "protection" sense.
10. The "diffuse catastrophe prior" is the exact analog of a Bayesian regression prior in which the underlying coefficient $\theta$ is an improper uniform distribution, and sampling then yields coefficient posterior means that are identical to those obtained in a classical regression (Zellner 1976). Thus, equivalent diffuse prior assumptions can be viewed as underpinning every reported classical regression analysis. Here too a different prior would yield different coefficient estimates and confidence intervals.
11. It makes no sense to generalize to a uniform continuous prior on the probability, because there needs to be a nonzero probability mass on zero disasters.
12. This "worst admissible magnitude" analysis is akin to the bounds in Hansen and Jagannathan (1991).

13. The expected compound simulated payoff is higher in the simulation than in the data because of compounding with resampled monthly equity premium draws. In this example, the investment would grow by 7.2%/year, to $7.80 over a hypothetical 323 months. The effect is analogous to the difference between the arithmetic and the geometric mean rate of return.
14. Another way to calculate the magnitude is to compare the expected rate of return on the put (the probability that a put will pay off in a disaster times the disaster's magnitude) with its monthly cost of rolling over.
15. In my view, good candidates are fortunate *ex post* coincidence (e.g., Jorion and Goetzmann 1999), declining risk premiums (e.g., Fama and French 2002), and increasing stock market participation and a fear of lower expected rates of return (e.g., Bansal and Yaron 2004).
16. The S&P 500 and the value-weighted stock return have an excess return correlation of 99.0%. It is almost impossible to distinguish between the two. Any differences are due to dividends having been paid.
17. Thus, if I had an option with 10 days to expiration but the futures were 100 days and 200 days to expiration, I would use a 70-days-to-expiration extrapolated futures price.
18. Within each area, the location of the mass can be indeterminate. The entire probability mass in the leftmost segment could be on one discrete outlier—say, $-\infty$.
19. The worst realization could be $-\infty$, also drawn with a 5% probability.

## References


Bansal, Ravi, and Amir Yaron. 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance*, vol. 59, no. 4 (August): 1481–1509.

Barro, Robert J. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, vol. 121, no. 3 (August): 823–866.

Chen, Hui, Winston Wei Dou, and Leonid Kogan. 2015. "Measuring the 'Dark Matter' in Asset Pricing Models." Working paper (30 June).

Coval, Joshua D., Jakub W. Jurek, and Erik Stafford. 2009. "Economic Catastrophe Bonds." *American Economic Review*, vol. 99, no. 3 (June): 628–666.

Fama, Eugene F., and Kenneth R. French. 2002. "The Equity Premium." *Journal of Finance*, vol. 57, no. 2 (April): 637–659.

Gao, George P., and Zhaogang Song. 2013. "Rare Disaster Concerns Everywhere." Working paper (October).

Gourio, François. 2012. "Disaster Risk and Business Cycles." *American Economic Review*, vol. 102, no. 6 (October): 2734–2766.

Hansen, Lars Peter, and Ravi Jagannathan. 1991. "Implications of Security Market Data for Models of Dynamic Economies." *Journal of Political Economy*, vol. 99, no. 2 (April): 225–262.

Jorion, Philippe, and William N. Goetzmann. 1999. "Global Stock Markets in the Twentieth Century." *Journal of Finance*, vol. 54, no. 3 (June): 953–980.

Kelly, Bryan, and Hao Jiang. 2014. "Tail Risk and Asset Prices." *Review of Financial Studies*, vol. 27, no. 10 (October): 2841–2871.

Kozeniauskas, Nicholas, Anna Orlik, and Laura Veldkamp. 2014. "Black Swans and the Many Shades of Uncertainty." Working paper, New York University and Federal Reserve Board (14 July).

Martin, Ian. 2011. "Simple Variance Swaps." NBER Working Paper 16884 (March).

Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics*, vol. 15, no. 2 (March): 145–161.

Nowotny, Michael C. 2011. "Disaster Begets Crisis: The Role of Contagion in Financial Markets." Working paper (7 March).

Orlik, Anna, and Laura Veldkamp. 2014. "Understanding Uncertainty Shocks and the Role of Black Swans." NBER Working Paper 20445 (August).

Rietz, Thomas A. 1988. "The Equity Risk Premium: A Solution." *Journal of Monetary Economics*, vol. 22, no. 1 (July): 117–131.

Rosenberg, Joshua V., and Robert F. Engle. 2002. "Empirical Pricing Kernels." *Journal of Financial Economics*, vol. 64, no. 3 (June): 341–372.

Seo, Sang Byung, and Jessica A. Wachter. 2015. "Option Prices in a Model with Stochastic Disaster Risk." Working paper, University of Pennsylvania (12 August).

Shiller, Robert. 2005. *Irrational Exuberance*, 2nd ed. Princeton, NJ: Princeton University Press.

Taleb, Nassim Nicholas. 2004. *Fooled by Randomness: The Hidden Role of Change in Life and in the Markets*. New York: Random House.

Wachter, Jessica. 2013. "Can Time-Varying Risk or Rare Disasters Explain Aggregate Stock Market Volatility?" *Journal of Finance*, vol. 68, no. 3 (June): 987–1035.

Zellner, Arnold. 1976. "Bayesian and Non-Bayesian Analysis of the Regression Model with Multivariate Student-*t* Error Terms." *Journal of the American Statistical Association*, vol. 71, no. 354 (June): 400–405.


 