CAPITAL MARKET BLIND SPOTS

JOSHUA COVAL, KEVIN PAN, AND ERIK STAFFORD*

October 2014

ABSTRACT:

This paper presents evidence suggesting that financial markets are susceptible to persistent mispricing in new financial insurance products. We argue that relatively sophisticated market participants often fail to notice that portfolios identified to be very safe by recently developed risk models may have state-dependent risk exposures arising from sources outside of the model. This error can create a perverse market dynamic that is sustained by market forces and academic research.

* Coval, Pan, and Stafford are at Harvard University. We thank John Campbell, Joshua Schwartzstein, Andrei Shleifer, and Jeremy Stein for valuable discussions and seminar participants at Harvard Business School, Indiana University, INSEAD, Stockholm School of Economics, NY Federal Reserve Bank, University of Michigan, and UT-Dallas.

Markets make mistakes. Mistaken notions of risk, and therefore pricing, occasionally persist in US capital markets for decades before becoming apparent to market participants. The realization typically arrives in the form of extreme losses to those specialized in bearing these risks. Once the realization occurs, pricing, capital requirements, issuance, and trading activity are significantly and permanently altered in the specialized segment of the capital market, suggesting that the market has indeed learned. A key question is whether these lessons are shared with other segments of the capital market or whether similar mistakes can persist elsewhere for long periods of time.

This paper investigates the possibility that capital markets develop blind spots when financial models are misapplied in real world capital markets. Our focus is on models that are relied upon by issuers of economic insurance products, such as index options and credit index derivatives. These are highly systematic exposures, for which there is likely to be natural demand. This creates incentives for relatively sophisticated market participants to develop low-cost methods for supplying these products. Any failure to recognize state-dependence in model errors will likely entail a replication shortfall in a poor economic state, or an abnormally low model-implied state price for poor economic states. A blind spot of this form results in the sale of economic insurance products that are underpriced and backed with insufficient capital.

We argue that a key reason why these blind spots emerge and persist for so long is that relatively sophisticated market participants rely upon learning rules (and research methodologies) that have little power to reject their model in economically benign states. Some participants may use the observation that market prices conform very closely to model prices as further evidence that their model is correct, gaining additional confidence in an incorrect view. Only when portfolios identified by the model to be very safe lose significantly more than believed possible is the model rejected. This learning event forces acknowledgment that the specialized model needs to be adjusted, but does not require a complete understanding of why. The updating process is complicated by the joint testing of the model and market efficiency. With a strong prior on market efficiency, the required *ad hoc* model adjustments are viewed simply as changes that are required to accommodate an evolving state-space, as opposed to evidence that the model was wrong all along and that imperfectly efficient market prices provided little power to detect this. This allows the same basic error to be repeated in other specialized areas of the capital market.

Merton (1997) recasts the logic of dynamic replication into a cost of production technology, whereby the original Black-Scholes (1973) and Merton (1973) option pricing models can be viewed as describing the cost of production of a derivative payoff in a perfect capital market. Replication shortfalls in poor economic states will contribute to the full cost of production an actuarial component from the shortfall; a risk premium due to the systematic nature of the shortfall; and a potential frictional cost of capital to the extent that covering the shortfall requires additional risk-bearing capital. All of these contributions strictly add to the cost of production implicit in a derivative model based on perfect capital markets, such that if derivative market prices conform closely to model prices, this potential shortfall is not being priced – or the possibility of state-contingent model errors is not being noticed. By failing to notice this possibility, derivative suppliers have a perceived cost of production that is lower than the true cost of production. Competition forces market price near the perceived cost of production, well below the true cost of production in this setting.

To the extent that the possibility of state-dependent model errors is explicitly recognized *ex ante*, these securities are likely to initially appear expensive relative to model prices *before* learning how well the model will perform in extreme economic states. Additionally, capital requirements are likely to initially appear to be quite conservative relative to those predicted by the model. In contrast to these predictions, we present evidence suggesting that suppliers of financial insurance products do not appear to charge much for this risk ex ante, suggesting that they grow highly confident in their ability to safely supply this insurance across market states, despite the fact that no adverse economic states have been realized over the life of the product.

The focus of our empirical analysis is on the stock market crash of October 1987 and the 2007-2009 financial crisis. We compare the performance of commonly-used strategies for assessing and hedging equity index put options before, during, and after the stock market crash of 1987 and study the ensuing repricing that occurred. We do the same on the credit side – examining the performance of risk models for corporate CDO tranches before, during, and after the financial crisis.

There are a number of striking similarities between these two episodes. In both instances, popular modeling breakthroughs (the Black-Scholes-Merton for index options and the Gaussian Copula for CDO tranches) coincided with dramatic rises in issuance and trading activity. The models performed remarkably well in the benign market conditions prior to the market corrections,[1] allowing risk managers and capital providers to gain confidence that the models were properly capturing risks. As a result, capital requirements for the suppliers of these exposures were small fractions of the notional value of the theoretically maximum payout. And, the market pricing of these securities seems to have included no additional risk premium to compensate suppliers for the potential risks associated with a large systematic shock, despite the fact that one had not yet been realized since the time that the innovation had been put into practice.

In both cases, large market corrections provided painful lessons to suppliers of these securities. A hedged writer of one-month at-the-money index options lost over 20% percent of the option strike price during the fall of 1987 – enough to easily exhaust funding requirements in place at the time. The seller of protection on the AAA-rated 7-15 tranche of investment grade CDX would have experienced losses of 35% percent from July 2007 through March 2009. These episodes resulted in broad and permanent repricing of the contracts in a manner consistent with a learning story – contracts with the greatest hedging errors were those whose insurance premia saw the largest permanent increases. Prior to the crash of 1987, out-of-the-money index options across all maturities exhibited little or no index skew: out-of-the-money puts had similar implied volatilities to those of at-the-money puts. Ever since the crash, out-of-the-money puts have had elevated implied volatilities relative to those of at-the-money puts. Whereas prior to the crash there was no difference, a 60 days to expiration, 50% out-of-the-money put has averaged a 30% higher implied volatility than that of an at-the-money put and has always been at least 20% higher.

During the five years that have followed the financial crisis, a similar repricing of CDO tranches has occurred. Prior to the crisis, for every 100 basis points of CDX yield, the 7-15 tranche averaged 2.7 basis points. Since 2010, this figure has averaged 14.3 basis points. The numbers are equally dramatic for the 15-100 tranche. Prior to the crisis, the senior-most 85 percent of the capital structure was promised 5.7 of every 100 basis points of CDX yield spread. Since 2010, this figure has averaged 29.3. Equally impressive has been how steady these shares have remained since the crisis

[1]We estimate that, prior to mid-October 1987, a modestly capitalized delta-hedged supplier of short-dated, out-of-the-money index put options would have never have experienced more than 3% annualized volatility on a broad inventory of options. Similarly, the writer of the 7-15 tranche written against a diversified index of 5-year North American investment grade credit default swaps (CDX.NA.IG) who delta-hedged by shorting the amount of the underlying index prescribed by the standard Gaussian Copula model would have never experienced a fully-funded draw-down greater than 1%.

– their standard deviations have been 1.7 and 4.1 basis points, respectively.

The fact that pricing permanently changed suggests that markets do, in fact, learn. However, the parallels between the evolution of the credit derivatives market around the 2007-2008 crisis and that of the equity derivatives market around the 1987 market crash are difficult to ignore – in both instances suppliers of insurance became overconfident in the ability of frictionless pricing models to help them manage risks that were highly systematic in nature – and suggest that the capital market has a difficult time generalizing its lessons and transferring them to other locations.

One key to validating the failing to notice story relative to other models of investor mistakes is that if one where to explicitly look for evidence of the neglected feature, they would find it.[2] This motivates some simple analysis of model errors prior to the major learning events. For example, we conduct a simple backtest to produce the daily hedging errors of a hypothetical index option supplier. We specify the relevant state variable to be daily SPX returns, and fit both a linear and quadratic function to the data. The linear specification suggests that the slope is precisely estimated close to zero. The quadratic specification confirms this result, but also shows that the quadratic term is reliably negative with a t-statistic of roughly -50, suggesting that if one had been looking for evidence of economically relevant state-contingent model errors they would have found it. A similar exercise for CDO pricing model errors also suggests that one looking for state-contingent model errors prior to the learning event would have found reliable evidence in that market too. Finally, we show that some common academic research methodologies suffer from the same weak power to reject the model against this particular alternative, effectively supporting the failing to notice mistake rather than offsetting it.

This paper proceeds as follows. Section 1 studies the pricing and risks of short-dated index options around the stock market crash of 1987. Section 2 studies the pricing and risks of CDX tranches around the 2007-2009 financial

[2] Our paper follows a long line of literature focused on learning in financial markets, including studies of learning about fund manager ability (Berk and Green, 2004, Jones and Shanken, 2005), learning about predictability in returns (Xia, 2001, Brav and Heaton, 2002, Cremers, 2002, Lewellen and Shanken, 2002, Brandt et al., 2005, Peng and Xiong, 2006, Garlappi et al., 2007, and Pastor and Stambaugh, 2009), learning and stock market bubbles (Pastor and Veronesi, 2003, Johnson, 2007), learning about one's investment ability (Gervais and Odean, 2001, Seru et al., 2010), and the consequences of learning for prices, volumes, and other observables (Harrison and Kreps, 1978, Romer, 1992, Welch, 1992, Timmerman, 1993, Wang, 1994, and Massa and Simenov, 2005) . See Pastor and Veronesi (2009) for a review. See Gennaioli, Shleifer, and Vishny (2012) for a recent model of neglected risk based on investors effectively truncating that state variable distribution.

crisis. Section 3 presents a simple model of a market for derivative securities. Section 4 discusses implications and Section 5 concludes.

1 Learning in the Market for Short-Dated Equity Index Put Options

We begin with an investigation of how the market learned about the risks of supplying index options around the stock market crash of October 1987. One way to think of this market and the closely-related portfolio insurance market is that there exists a natural supply of investors who would like to insure against large stock market losses creating an incentive for low cost producers to supply insurance contracts to meet the demand. From the perspective of an insurance contract, the policy supplier would typically set aside capital (reserves) to cover possible losses. Since the risks that these policies are insuring are essentially perfectly correlated, diversification offers little benefit in reducing the required amount of capital as the number of policies increases, making this a highly capital intensive business. In 1973, Black-Scholes and Merton (BSM) developed option pricing models that describe a lower cost production technology – dynamic replication. In a perfect capital market with diffusive underlying price shocks, option payoffs can be precisely manufactured with a riskfree production technology, requiring only a small amount of startup capital, specifically the initial option value. The replicating portfolio cumulates profits and losses to exactly match the intrinsic value of the option at each point in time and precisely delivers the terminal contract payoff at option expiration. Equivalently, an option can be perfectly hedged by shorting the replicating strategy.

In 1983, the Chicago Board Options Exchange introduced trading in index options. From the start, the market experienced explosive growth. Within one year, trading volume exceeeded 77 million contracts and by 1986 trading volume exceeded that of all individual equity options. At this time, many of the longer-dated index options were sold in the form of an over-the-counter product known as *portfolio insurance*. This product, pioneered by the firm Leland, O'Brien, Rubenstein (LOR), would use the Black-Scholes hedging prescription to replicate the outcomes of a three-year, at-the-money put option on the S&P 500 index. Along with trading volume on the exchange, portfolio insurance experienced explosive growth. By 1987, approximately three percent of all equity assets were covered by different varieties of portfolio insurance, the majority of which was on behalf of conservative pension funds.

The explosive growth of the index options and portfolio insurance markets was driven in no small part by the success of the Black-Scholes model in explaining its prices and risks. The frictions of real world capital markets appeared to create only minor distortions relative to the predictions of the model based on frictionless markets. The market pricing of index options was highly consistent with model prices implied by estimated model inputs, and importantly the risks that insurance suppliers were bearing seemed to be well hedged by those following the model's prescribed delta-hedging strategy.

To see this, consider the experience of an index put option supplier (insurer of stock market losses) relying on the dynamic replication strategy. Each day, we calculate the hedging errors across a portfolio of short-dated out-of-the-money index put options that are hedged according to the BSM model. Each day the option supplier holds an option portfolio consisting of one short put option at each moneyness level ranging from 0.5 to 1.0, with an increment of 0.025 (21 options) for each of the closest 3 monthly maturities, holding a total of 63 options. We calculate hedging errors simply as:

$$HE_{i,t+1} = modelDelta_{i,t} \cdot \Delta S\&P_{t+1} - \Delta Option_{i,t+1} \tag{1}$$

where the model delta is calculated using an estimate of market volatility. Specifically, each day's volatility estimate is based on a GARCH(1,1) model using daily SPX price changes from January 1926 to the prior trading day. We estimate the market implied volatility each day as the GARCH 21-day forecast volatility multiplied by 1.18 to match the mean of the VIX where both time series are available (see the Appendix for a comparison of the GARCH volatility estimate to the CBOE Volatility Index (VIX)). To meaningfully compare these dollar hedging errors across contracts and through time, we calculate daily hedged returns to the option supplier. We assume that the option supplier contributes 100% of the strike price of each option as collateral against these short positions, such that these positions can be viewed as being *unlevered*. The daily returns to the hedged option portfolio are calculated as:

$$R_{p,t+1} = \sum_i HE_{i,t+1}/Capital_t. \tag{2}$$

Note that this calculation is isolating the costs of production and excludes any net profits from capturing bid-ask spreads (i.e. market making), which were quite large during this period.

Figure 1 displays the daily portfolio returns for this hypothetical equity index put option supplier from January 1927 through June 2012. Over the period from 1973, when the Black-Scholes-Merton option pricing model was published, to September 1987, the option supplier's hedged portfolio (unlevered) was nearly riskfree with annualized volatility of 0.83%[3], suggesting that many option suppliers would have had favorable experiences with this strategy for several years prior to mid-October 1987.[4] Market pricing of these contracts was largely consistent with BSM model because market participants believed that the model accurately described the risks of supplying these contracts, and the realized experiences of participants offered nothing to challenge these beliefs for a long period of time.

On October 19, 1987, these beliefs were severely challenged, when the hypothetical unlevered option supplier with a long history of producing essentially riskfree returns lost over 12% in one day. Given the volatility of the strategy up to this point, this return represented a standardized shock (Z-score) of -240. The suppliers of short-dated equity index put options quickly realized that their technology was, in fact, not riskfree when the S&P 500 dropped 30% over four consecutive trading days, down 20% on October 19, 1987. Moreover, these losses were not just random hedging errors that were easily diversified across a large portfolio of risks. These losses coincided with a 30% drop in the US stock market! As Peter Bernstein put it in *Against the Gods: The Remarkable Story of Risk*, "The cost of portfolio insurance in that feverish market turned out to be much higher than paper calculations predicted."

1.1 The Market's Pre-Event Information & Belief Set

The dynamic replication technology for supplying out-of-the-money index put options suggests that the amount of capital required to reliably deliver each contract's terminal liability value is a small fraction of the strike price (i.e. the theoretically maximum value). For example, with annual stock market volatility of 20%, the maximum of each BSM option price divided by its strike price is less than 4%, implying that leverage of 25x would be highly conservative since most of the contracts require far less capital. Of course, at this leverage level, the hypothetical option supplier would have

[3] This same calculation using actual option prices from the Berkeley Options Database over the period January 1986 through September 1987 produces an annualized standard deviation of 0.85%.

[4] Index futures began trading on the Chicago Mercantile Exchange in April 1982, and index options began trading on the Chicago Board Options Exchange in 1983.

experienced a daily return worse than -300%. In reality, before October 1987, the Chicago Board Options Exchange (CBOE) required: "5% of the aggregate index value, minus any out-of-the-money amount plus 100% of the current premium, to a minimum of 2% of the current aggregate index value plus the current premium."

Table 1 tabulates the CBOE capital requirements for various one-month to maturity index put options, along with their BSM model prices, assuming the current level of the index is 100, a riskfree rate of interest of 5%, a dividend yield of 2%, and annual volatility of 20%. The capital requirements in place at the time of the event are slightly higher than those implied literally by the BSM model, but are very small relative to where they eventually end up and relative to the losses that an unlevered hedged option supplier experienced on October 19, 1987.

From a practitioner's perspective, the period from 1973 through September 1987, nearly 14 years, represents a significant portion of one's career. Moreover, such a career would have likely been hugely successful. From the perspective of a longer time period that includes the 1930s, the October 1987 event does not look too extreme. Figure 1 shows that the hypothetical unlevered option supplier would have experienced many daily losses exceeding 6%, and that these losses would have coincided with generally awful economic conditions.

Figure 2 characterizes the risks of the underlying equity index in two ways. The first panel displays the daily drawdown of the market, calculated each day as the current index level relative to the previous maximum level of the index, as a percentage. The notable option supplier losses from Figure 1 clearly occur at times when the stock market index is well below its previous high, indicating that recent returns on stocks have been poor. From the perspective of drawdowns, there is nothing remarkable about the October 1987 event, as the market is down from its highs by more than this amount quite often in the historical period. The second panel displays daily return Z-scores, scaled by the GARCH(1,1) volatility estimates. The daily Z-scores reveal that the uniqueness of the October 1987 event is its speed of drawdown, as this is the worst one-day return shock in the sample at -12. However, it is only slightly worse than the previous worst shock of -10.

Taken together, it appears that market participants discounted the historical data very strongly, basing their belief on the much shorter realized experience with the technology (1983-1987). Over this shorter sample, estimates of the total risk of the strategy suggest it is nearly riskfree and that there is no evidence of the risks being systematic. The confidence in these views – producing allowable leverage levels and market pricing so consistent with the model prices – would not be permissible without essentially ignoring the longer time series.

1.2 The Market's Lesson

The underlying demand for insurance against large stock market losses continued after this episode. The technology for supplying this insurance needed to be improved. The essence of the modifications to the technology were to continue to hedge the diffusive risks (including the additional sources of risk that arise from stochastic volatility), but to also contribute additional risk-bearing capital to cover the large hedging errors caused by the possibility of price changes that were not well described as diffusive shocks. Equivalently, leverage was reduced considerably. Immediately after "Black Monday," on October 26, 1987, the CBOE essentially doubled capital requirements, and then doubled them again on May 19, 1988 (see Table 1). Importantly, unlike in the BSM model, the additional required capital is now recognized as risk-bearing capital, which, in addition to occasionally being drawn upon to cover the losses associated with large hedging errors, will require a large risk premium given the highly systematic nature of these large losses. These two features combine to create a significant wedge between market prices and the original BSM model.

The pricing of equity index options has been permanently different ever since this event. Most noticeably, there is a pronounced implied volatility skew, whereby further out-of-the money options are progressively more expensive than the BSM model predicts. Figures 3a and 3b plot the daily estimated implied volatility skew (slope of the implied volatility as a function of moneyness) before and after the crash using option data from the Berkeley Options Database. Each day, t, we calculate Black-Scholes/Merton implied volatilities for each put option, $IV(t, DTM, K)$, where the strike price, K, is less than or equal to the prevailing index level, S_t, excluding strike prices more than 2 standard deviations ($VIX_t \cdot \sqrt{\frac{DTM}{365}}$) below S_t. We estimate the daily IV skew for a given DTM from an OLS regression of IV on K, requiring at least five IV observations. There is quite clearly a structural break in the pricing of index options around the October 1987 stock market crash. Moreover, this pattern is highly consistent with the magnitude of the realized hedging errors on October 19, 1987, scaled by the prevailing market prices. The hedging errors were expected to be small relative to the model price, which, in theory, represented the total amount of capital required at each point in time to supply the terminal option payoff. To the extent that the hedging errors are large relative to the model's implied capital requirement, additional risk-bearing capital will be required. Index option prices were repriced in proportion to their percentage hedging errors. Figure 4 displays the log of these percentage errors for each of the holdings in the hypothetical option supplier's portfolio on October 19, 1987, matching the well-known pattern in BSM model implied volatilities – in both moneyness and time-to-maturity – that has persisted ever since then.

A final consequence of the increased cost of supplying economic insurance in the options markets was that less demand was interested in being fulfilled at these higher prices. Index option trading volume took almost twenty years to revisit the 80 million contracts level achieved in 1986. And portfolio insurance all but died out within a few years.

2 Learning about the Risks of Securitizations

Twenty years after the introduction of index options and portfolio insurance, the market for Collateralized Debt Obligations (CDOs) began to take off. Following several years of modest growth, between 2003 and 2006, global annual CDO issuance increased exponentially, from under $80 Billion to over $450 Billion. The attraction of the CDO, and securitization more generally, was that it was a means of manufacturing safe assets by pooling economically-linked cash flows into diversified portfolios and partitioning these cash flows into prioritized claims called tranches. The risks of the underlying asset pool were generally viewed to be confined to the junior tranches, implying that the often "over-collateralized" senior tranches would be very safe.

Once again, the catalyst for the explosive growth appears to have been the development and adoption of new financial models that convinced investors and rating agencies that the risks of supplying economic insurance – i.e. the risks in the senior tranches – were minimal and could be easily managed. The key modeling breakthrough appears to have been the Gaussian Copula model of Li (2000). Within a short period of time, the Gaussian Copula and similar models were widely adopted by structurers of CDOs and their risk managers. In August of 2004 Moody's put the Gaussian Copula at the heart of their CDO ratings methodology. Within a week, S&P did the same.

Market participants rapidly gained confidence in the models' characterization of senior tranches as extremely safe. Evidence that they were widely viewed to be safe comes from their pre-crisis pricing, their capital requirements, the haircuts required when they were used as collateral in repo transactions, and the fact that US Broker-Dealers appear to have relied so heavily on non-Govt securitizations as collateral in repo transactions.

To explore how they gained confidence so quickly that the "safe assets" created through securitization were indeed safe, we again examine the historical experience of hypothetical market participants. We focus on the Dow Jones CDX North American Investment Grade Index (CDX) and its tranches. This index is an equally-weighted portfolio of 125 liquid 5-year credit default swaps on investment grade (IG) bonds. The CDX is rebalanced every March and September to reflect changes in the composition of the liquid investment grade corporate bond universe, notably dropping securities that are deemed to have sufficiently increased credit risk. Our data cover the period September 2004 through September 2012, and include spreads on the CDX and the spreads on the CDX tranches. CDO tranches are derivative securities with payoffs based on the losses of the underlying CDO portfolio. The CDX tranches are defined in terms of their loss attachment points. For example, a $1 investment in the 7-10 tranche receives $1 if the total losses on the CDX are less than 7%; $0 if the total losses on the CDX exceed 10%; and a payoff that is linearly adjusted for CDX portfolio losses between 7% and 10%. The structure of the CDX tranches have changed over our sample period. For example, in September 2010, the 7-10 and 10-15 CDX tranches were combined into a single 7-15 tranche; the 15-30 and 30-100 tranches were merged into a single 15-100 tranche.

For ease of exposition, we will primarily focus on wider tranches, as these hypothetical merged tranche spreads can be constructed historically, while the hypothetical separate tranche spreads cannot be determined without a specific model. Additionally, the market conventions for quoting CDX tranche prices has changed considerably through time in terms of upfront and running spreads. Again, for expositional convenience we will sometimes focus on all-in upfront prices of protection when describing credit spreads, and subtract these insurance premia from 5-year zero coupon riskfree bonds when describing security prices.

Prior to the financial crisis, the 7-10 CDX.IG tranche was widely considered to be a AAA-rated credit security by market participants. Consequently, a 7-100 CDX.IG tranche would be viewed as even safer. Over the period January 2005 through June 2007, the credit spread of the 7-100 CDX tranche averaged only 0.03%, suggesting that market participants viewed this to be a highly safe security. This economically benign time period was generally associated with low credit spreads. The CDX index, essentially reflecting the investment grade corporate sector, had an average credit spread of only 39 basis points over this period. The reason that the 7-100 tranche could be so safe was because of the extreme confidence that the risks were confined to the junior 0-7 tranche, which had an average credit spread of 4.4% over this period.

Market participants specializing in credit risk are generally viewed to be highly sophisticated (see Collin-Dufresne, Goldstein, and Yang (2012)) with access to liquid hedging strategies that allow them to be low cost suppliers of derivatives like CDS and the CDX tranches described above. To describe the experience of a CDO underwriter relying on a hedging strategy to reduce the risk of its safe tranche inventory, we calculate the daily hedging errors for the 7-15 and the 15-100 CDX tranches, hedged with their empirical sensitivity to the underlying CDX. We calculate hedging errors simply as:

$$HE_{i,t+1} = \widehat{\Delta}_{i,t} \cdot chgCDX_{t+1} - chgTranche_{i,t+1} \tag{3}$$

where the empirical delta, $\widehat{\Delta}_{i,t}$, is calculated using changes in the all-in upfront prices of protection over the past 100 trading days. To meaningfully compare these dollar hedging errors across securities and through time, we calculate daily hedged returns to the underwriter. We assume that the tranche underwriter contributes capital of $0.20 per $1 of notional tranche exposure for the 7-15 tranche and $0.08 per $1 notional for the 15-100 tranche. These capital contributions are considerably more than required at the Federal Reserve discount window for AAA-rated securitizations, which were 4% at the time.[5] The daily returns to the hedged tranche are calculated as:

$$R_{i,t+1} = HE_{i,t+1}/Capital_{i,t}. \tag{4}$$

Note that this calculation is isolating the costs of production and excludes any net profits from underwriting fees, which were quite large during this period. In fact, CDO issuance was high and growing rapidly over this period, generating both fees and increased underwriter inventory of safe tranches.

Panel A of Figure 6 displays the cumulative daily hedged returns for this hypothetical tranche underwriter from September 2005 through December 2012. Panel B plots the daily hedged returns over this period. Over the 20 months prior to July 2007, the tranche underwriter's hedged tranche positions were nearly riskfree, each with annualized volatility of about 2.5%. In fact, the capital supporting each of these positions was roughly 100x the

[5] Since these securities are viewed to be very safe (e.g. AAA-rated) one could argue that the delta should be zero. The analysis produces qualitatively similar results with delta set to zero.

largest absolute value 1-day dollar hedging error realized over the pre-crisis period.

Beginning in July 2007, the hypothetical underwriter's experience changed dramatically. The annualized volatility of the July 2007 daily hedging errors were 37% for the 7-15 tranche and 15% for the 15-100 tranche. Moreover, they were on average negative, producing large losses for the underwriter. As the figure makes clear, things would only get worse for the underwriter over the next 18 months as the hedging strategy completely fails to eliminate the risks of these securities that were believed to be safe by so many market participants.

Since the financial crisis of 2007-2008, the pricing of securitizations has been permanently different. The previously held view that the risks of the underlying pool were confined to the junior tranche has been significantly updated. Figure 7 displays the time series of the market's assessment of where the risks of the CDX lie. In particular, the figure simply plots the all-in price of protection for $1 of notional for each tranche, weighted by each tranche's share of the capital structure, offering a model-free perspective of the market's beliefs through time. During the pre-crisis period prior to July 2007, while the hypothetical tranche underwriter's hedged returns were essentially riskfree, the safe portion of the CDX capital structure (the 7-15 and 15-100 tranches) claimed only 10% of the total risk, with the junior 0-7 tranche claiming 90%. During the crisis, and ever since, the safe portion of the capital structure claims nearly 50% of the risk of the underlying asset pool. Put differently, the market now views the senior tranches to be five times as risky as it did prior to the crisis.

Consistent with the market's updated view that the "safe" securities created through securitization are not that safe, there is enormous effort re-evaluating how financial institutions should rely on non-Government securitization-based repo for their funding, and there has been essentially no issuance of these securities in the 4 years following the crisis (SIFMA), as illustrated in Figure 8.

3 How Markets Learn

The two episodes described above – which arguably represent the two largest learning events in recent financial market experience – share a number of common features. In both cases, underlying the rise in each market was an excess demand for products that isolated and, through hedging or over-collateralization, eliminated exposure to systematic economic risks. In both cases, a key modeling breakthrough helped guide suppliers of these products in allocating capital, hedging, and pricing. This allowed new securities to be issued and traded in a new, specialized segment of the capital market. As the securities trade, evidence builds that the model is effective and participants gain confidence that it is properly assessing risks, including and particularly in states that have not yet occurred. A learning event then takes place, wherein securities and portfolios identified by the model as safe experience significant losses. This triggers increases in capital requirements, a broad and permanent repricing of securities, and fundamental adjustments to the model.

The key challenges for a story of this dynamic to explain are (1) how does confidence build, given the widely accepted view that all models are false and that it is the relatively sophisticated market participants that will end up on the wrong side of the learning event; and (2) what limits the lesson of the learning event to the specialized market without being shared generally, thereby allowing this pattern to repeat? Our contention is that relatively sophisticated market participants – i.e those who rely on models – make a single but significant mistake in their assessment of the model: *They fail to notice the possibility of state-contingent model errors when supplying economic insurance, and therefore, they fail to anticipate and appreciate the resulting consequences for equilibrium.*

This is a first-order mistake. Since all models are false they will produce errors – hedging errors and pricing errors. Evidence of state-contingent model errors would likely invalidate the model. But to the extent that market prices conform closely to model prices – perhaps because many influential participants are relying on the same model – this will likely lead to statistical tests failing to reject the joint hypothesis of the model being correct and this segment of the capital market being efficient. Market participants will gain confidence that the model is correct and the market is efficient when in fact the opposite is the case.

Moreover, participants that ignore the possibility of state-contingent model errors will become the low cost providers of economic insurance and therefore highly influential in determining market prices. In the extreme, the prices offered will literally be those implied by the model, which will become the data used in statistical tests to validate the model. Clearly, this represents a weak test against the alternative that both model and market prices are wrong.

3.1 Derivative Pricing through Replication

Suppose there is a specialized segment of the capital market subject to the mistake described above. Specifically, we assume that priors are updated through rules and testing methodolgies based on whether model errors are "small," which is a necessary property of model errors, but insufficient for ruling out the possibility that large errors will emerge in bad states.

The logic developed here adopts the "cost of production" terminology used in Merton (1997). Consider a derivative whose payoff, $f(s_T, X)$, can be expressed as a function of the underlying security's price at expiration, s_T as well as the theoretically-maximum possible payout, X. Assume complete agreement among all market participants about the perfect capital market (PCM) value of the derivative payoff, $P^{PCM} = g[f(s_T, X)]$, where $g[]$ represents the model-implied perfect capital market cost of production. This model is taken to the real world, where derivative suppliers are required to contribute a minimum amount of risk-bearing capital to support their exposure. In particular, define θ as the fraction of the maximum loss that is required to be held by the derivative supplier per contract, so that risk-bearing capital is simply the supplier's total exposure multiplied by θ,

$$k = \theta \cdot X \cdot q,$$

where q is the quantity of contracts sold. Suppose a simple rule is used for the determination of required capital – suppliers are required to post the *unlevered* state-contingent hedging error at some critical state variable realization, α^*, (i.e. stress test):

$$\theta = -\boldsymbol{\varepsilon}(\alpha^*). \tag{5}$$

Finally, if we assume a *frictional* cost of risk-bearing capital, $C(k)$, that is increasing and convex[6] in the quantity of risk-bearing capital, $C(0) = 0$, $\frac{\partial C}{\partial k} > 0$, $\frac{\partial^2 C}{\partial k^2} > 0$, the full cost of production per unit can be described as:

$$G = g + a + b + \frac{C(q \cdot \theta \cdot X)}{q}. \tag{6}$$

where a is the actuarial contribution of real world hedging errors (mean cumulative hedging error), b is the risk premium required for any dependence

[6] See Froot, Scharfstein, and Stein (1993) for a discussion of the nature of external financing costs and their implications for firm risk management activities.

between real world hedging errors and relevant state variables, and C is the contribution from the frictional cost of capital.

The effectiveness of hedging affects prices, and is encapsulated in the function that connects model hedging errors to state variable realizations, $\boldsymbol{\varepsilon}(\alpha)$. If hedging errors are viewed to be small across all states, both the supplier's subjective valuation of these errors and the required risk bearing capital will be small, leading to a minimal deviation from the frictionless market model price. On the other hand, if these are viewed to have a strong negative relation, such that a hedged-portfolio is expected to experience significant losses in bad economic states, then there can be a significant wedge between the full cost of production and the one implied by the frictionless model.

The key to the "failing to notice" mistake is that the structure of the problem is not properly understood. Our setting of economic insurance products ensures that the ignored feature of the problem is highly relevant. Consequently, updating rules that ignore this feature of the problem have the potential to be meaningfully misleading, regardless of how priors are initially formed.

It is plausible that the model itself may anchor priors, since market participants did not have a compelling way to think about the problem without the model. Financial pricing models that can be solved typically assume perfect capital markets, and therefore represent a *lower bound* on the full cost of production for economic insurance products, $G \geq g$. For example, the perfect capital market BSM option pricing model implies no hedging errors ever, such that $\varepsilon(\alpha) = 0$ across all states. This implies that $a = 0$, $b = 0$, $\theta = 0$, $C(0) = 0$, and therefore that the full cost of production is simply the frictionless market cost of production, or the BSM model price. The model does not explicitly suggest how it should be generalized. The key insight of analyzing $\varepsilon(\alpha)$ is not explicitly suggested by the model. This needs to be learned.

Failing to notice amounts to not imposing the proper structure on the conditional hedging error distribution. Without imposing the proper structure, one is left with relatively weak empirical methods to recover this distribution. For example, one might estimate value-at-risk models based on distributional assumptions, or measure sample volatility of hedging errors and their correlation with relevant state variables, which can be used to linearly forecast what would happen in an unobserved bad economic state.

While all models are false, many models can produce small hedging errors with low correlation to economic state variables in economically benign times. The hedging error analysis will be highly consistent with the predictions of the frictionless market model, but also with many other models. The sample will provide no evidence against the specific alternative of signficant state-dependence in bad economic states, $\varepsilon(\alpha^*) < 0$. These two highly different views of capital markets may not be distinguished by the empirical evidence, so confidence that one of them dominates is inappropriate.

Failing to notice explains why this alternative is not considered, and therefore how confidence in the wrong view can cumulate. With many high frequency observations suggesting that hedging errors are small and mostly uncorrelated with relevant economic state variables, there is no need for much risk-bearing capital. In particular, by analyzing hedging errors generated in economically benign environments with statistical methods that do not explicitly look for dependence with bad economic states, one can incorrectly form the view that $\varepsilon(\alpha^*) \approx 0 \Longrightarrow \theta \approx 0 \Longrightarrow C(\approx 0) \approx 0 \Longrightarrow G \approx g$. The perceived (but wrong) full costs of production will be those associated with the frictionless capital market model.

From here, market forces take over, with competition enforcing the condition that prices equal their perceived marginal costs of production, $P^{Mkt} \approx g$. Empirical studies will confirm that $E[e] \approx 0$, allowing researchers to conclude that the joint test of market efficiency and the model is not rejected.

Interestingly, this may also lead some participants to improperly conclude that frictions do not matter. This inference error could improperly reinforce other frictionless market beliefs like rational expectations and limit one's openness to models that highlight the role of frictions. A blind spot of this sort may limit the transferability of the lesson to other areas of the capital market where market frictions may be important, allowing the failing to notice mistake to be perpetuated.

The fundamental mistake of "failing to notice" the state-dependence of model errors combines with normal market forces to produce a fascinating market dynamic. The notion of failing to notice, is inspired by the "learning through noticing" model of Hanna, Mullainathan, and Schwartzstein (2014), who apply their model to a sample of experienced seaweed farmers. They find that relevant input dimensions of the farmer's production technology are ignored and show that the resulting input choices are far from optimized in these dimensions. Moreover, they find that learning does not occur without the farmers receiving explicit summaries of how the previously ignored dimensions affect outcomes.

Failing to notice the proper structure of the situation is also central to the dynamics of the self-fulfilling prophesy (Robert. K. Merton, 1948).

> "The self-fulfilling prophesy is, in the beginning, a *false* definition of the situation evoking a new behavior, which makes the originally false conception come *true*. The specious validity of the self-fulfilling prophesy perpetuates a reign of error. For the prophet will cite the actual course of events as proof that he was right from the very beginning."

This dynamic appears to not be recognized by the relatively sophisticated market participants. Their frictionless market model is surely wrong, but seems to work. Because it works, it is used more and more. As it is used more, market prices conform to the model. Tests of the model confirm this, establishing that the model is correct and that markets are efficient. Failing to notice allows for the market forces channel to create this self-fulfilling feature, and for market participants to use empirical evidence that has little power against a specific alternative (i.e. painful state-contingent hedging errors) to update their views of the situation.

It seems unlikely that confidence could build quickly around a frictionless model if market participants were actively looking for evidence of state-contingent hedging errors, an idea we explore in the next section. It seems more likely that market participants fail to notice this dimension of the problem, like the seaweed farmers. They simply are not looking for it, it is not hitting them over the head, and they are otherwise satisfied with their understanding of the situation. What is remarkable here is that these are the relatively sophisticated market participants and not seaweed farmers.

To understand how the affected specialized market learns, but in a way that allows the lesson to not become generally known, we investigate the interpretation of research around these events.

3.2 Noticing State-Contingent Model Errors

Consider the estimation problem before the new technology has been used, and therefore before any data exist. This is a hard problem, even assuming that the α distribution is known. Neither the functional form nor the parameters are known.

What if one was to look for state-contingent model errors? One approach would be to conduct a backtest of the hedging strategy, as was done in the previous section, and analyze the resulting hedging error distribution to form a prior distribution.

3.2.1 Short-Dated Index Options

Taking this exercise seriously, we conduct a simple backtest of the hedging errors of the hypothetical index option supplier described in Section 1. For daily hedging errors, we specify the relevant state variable (α) to be daily SPX returns, and fit two simple functions to the data prior to the introduction of index options in 1983 (linear and quadratic), displayed in Figure 9 and in Table 2. Both specifications suggest that hedging errors are typically small, averaging just a few basis points per day. However, this misses the point. The important question is what will happen in a bad economic state. Each specification can be used to determine a required capital level, $\theta = -\varepsilon(\alpha^*)$. To evaluate, we assume the critical state variable value to be covered by the supplier is $\alpha^* = -20\%$. The quadratic specification produces a required capital of 10% that is materially higher than the linear specification, which suggests 0%.

Interestingly, the time series evolution of these estimates is quite stable, either including or excluding the pre-1983 backtest data. In every subsample reported, the quadratic term is negative and highly statistically significant, suggesting that losses on the hedged portfolio will be reliably larger than the 2% capital requirement in place at the time.

The implications of this analysis are straightforward, and quite at odds with the evolution of the capital markets we examined earlier. The data clearly suggest a pronounced nonlinearity in the state-contingent hedging errors that would lead someone forming a prior belief about initial capital requirements and the valuation of cumulative hedging errors far above where actual market participants did. Moreover, even if market participants had not back-tested to form their initial priors, by September 1987, a market participant actively looking for an indication of state-contingent hedging errors would have found reliable statistical evidence.

Accounting for Model Uncertainty: A growing literature emphasizes the role of model uncertainty in forming and updating beliefs about market outcomes. In our setup, when $\varepsilon(\alpha)$ is unknown and must be estimated, model uncertainty is likely to add to the full cost of production. Given the nature of the problem – realizing hedging losses in bad economic states – it is likely that one will not only be inclined to go with the specification that produces the highest required capital, but to require an additional cushion to the estimated θ, which will increase costs of production. Similarly, the suppliers' assessment of the valuation of hedging errors will be higher in the presence of model uncertainty. The consequence of model uncertainty –

for derivative prices offering economic insurance – will be for prices to start quite high and then decline as model uncertainty is resolved and for capital requirements to start higher than their long-run equilibrium level. This is clearly inconsistent with the actual evolution of these markets, where capital requirements approach their eventual level from below.

3.2.2 Investment Grade CDX Tranches

In many ways, this exercise is more interesting for the investment grade CDX tranches, as market participants have the benefit of nearly 20 years to process the events of 1987, and determine their relevance for a new model-driven innovation.

There is reason to expect *ex ante* that the payoffs of derivatives on large diversified portfolios of corporate credit securities will be related to the performance of the S&P 500 index, so we maintain this particular state variable (see for example, Coval, Jurek, and Stafford (2009) and Collin-Dufresne, Goldstein, and Yang (2012)). Consider the [7, 15] tranche of the 5-yr CDX.IG. A structural model mapping the systematic risks of corporate assets into equity index options would suggest that this tranche resembles a portfolio of US Treasury bonds and a short position in an equity index put spread. In particular, we focus on a simple equity index put spread that writes q 5-yr equity index options struck at $K_U = 50\%$ moneyness and buys q 5-yr equity index options at $K_L = 30\%$ moneyness, where $q = \frac{1}{K_U - K_L} = 5$. This tranche was widely considered to be safer than the underlying investment grade portfolio, with market participants viewing it to be equivalent to AAA-rated, and its average credit spread was lower than that of the CDX. As such, even unhedged, this portfolio was identified by popular financial models as being extremely safe, requiring minimal risk-bearing capital.

To backtest the performance of how this portfolio might be expected to perform across a wide variety of economic states, we repeat the basic exercise above based on a statistical procedure to map GARCH(1,1) into implied volatility. In particular, we use the simple GARCH(1,1) volatility model to forecast 5-yr expected volatility and then regress this forecast onto the corresponding daily implied volatility for 5-yr at-the-money equity index options over the period 2003 through 2012. We conservatively assume that there is no implied volatility skew in 5-yr implied volatilities. Since the GARCH(1,1) estimates are available back to 1926, we have a long time series of hypothetical tranche prices to examine. Figure 10 displays the actual [7, 15] tranche value across various levels of SPX drawdowns over the period 2004-2012. The figure also displays the values of the equity index put option replicating portfolio $[K_U = 50\%, K_L = 30\%]$ based on actual option prices and the GARCH-based backtest option prices. All three plots look remarkably similar – all suggesting a strong systematic risk exposure over this period.

However, our interest is in understanding what was knowable at the beginning of the sample. What does the backtest suggest before this period. Figure 11 displays this results from this analysis conducted over the period 1926 to 2000, along with a version for the underlying investment grade CDX, which is modeled simply as a portfolio of US Treasury bonds and a single short equity index option at 60% moneyness. Two things are striking about the analysis. First, a very strong systematic risk profile exists in the pre-2000 backtest. Second, in both the pre-2000 backtest and the actual data, the 7-15 tranche has a much steeper systematic risk profile than the CDX in bad economic states, suggesting that it will lose *more* of its value in bad states than the CDX, despite the market's pre-crisis assessment that it is safer.

3.3 How Research can Fail to Notice

Prior to 1987, the Black-Scholes pricing model was hailed as a great success. A number of studies documented the impressive degree to which option prices conformed to the Black-Scholes formula. Perhaps the most thorough empirical study of options was carried out by Rubenstein (1985), in which he revealed that equity options trading on the CBOE from August 1976 through August 1978 were highly consistent with one of the central predictions of the Black-Scholes formula – that options with different strike prices or maturities should have the same implied volatility. As Ross (1987) put it, "When judged by its ability to explain the empirical data, option-pricing theory is the most successful theory not only in finance, but in all of economics."

3.3.1 Looking for Small Errors Instead of State-Dependent Errors

To see just how well the model priced index options prior to the October 1987, consider an evaluation of SPX index options using the approach of Rubenstein (1985). On each day, the volatility is identified that minimizes RMSE for all put options that traded on that day. Figure 12 presents the median RMSE for each month from 1983 onward. As can be seen, the single-volatility parameter is sufficient to price index options almost perfectly prior to the crash of 1987. On over half of the days, all traded options can be priced to within a \$0.50 bid-ask spread. This means that the implied volatility surface was perfectly flat across strike prices and across maturities on most days prior to October 1987. More importantly, it means that option sellers were not charging any additional premium to cover any possible frictional costs of production. Given the proper relation implied by competition, $P^{Mkt} = G$, the joint test of market efficiency and a frictionless capital market model should produce biased pricing errors equal to $E\left[P^{mkt} - g\right] = E\left[a + b + C(k)\right] > 0$. This pattern is clearly illustrated in Figure 12.

While it may have been tempting to conclude that Black-Scholes was indeed correct during this period, it's worth remembering that any test of an asset pricing model is actually a joint test of the model and market efficiency. Black-Scholes may have been perfectly consistent with market prices prior to 1987, but that does not imply that the market was necessarily assigning correct prices to index options. And, to the extent that a popular arbitrage strategy implemented by option traders at the time was a form of spread trade identified by deviations from the Black-Sholes formula (see MacKenzie, 2006, p. 164), this would work to equalize implied volatilities across strike prices. In this way, the data being used to test the model were those being produced by the model. The tests were only revealing what model was being used by those setting the prices – not that they were necessarily using the correct model or setting correct prices. Adding that to the fact that index options had only been trading for only a few years – a period with relatively modest realized volatility – suggests that only somone failing to notice the possibility of state-contingent model errors should be eager to declare victory at that point.

3.3.2 Adding a Jump Process to Explain Prices *Ex Post*

As we discussed above, October 1987, brought about dramatic and permanent changes in index option prices. Given the large increase in out-of-the-money put options, the Black-Scholes formula with a single volatility parameter could no longer be considered an accurate description of index option prices. As Figure 12 shows, the RMSE from the Black-Scholes model goes from being negligable to averaging around 10 percent.

However, it turns out that a relatively simple modification can capture a reasonable amount of the change in pricing that occurred. In particular, by adding a systematic jump process to the log-normal diffusion at the heart of the Black-Scholes formula, the model can now assign relatively high prices to options most exposed to a large decline in the index. Specifically, consider adding a jump process where the jump represents a 50 percent decline in the index and then identifying both the risk-neutral volatility and jump intensity that best describe option prices.

As we see in Figure 12, this approach continues to price options almost perfectly prior to the crash of 1987, and it eliminates a good fraction of the Black-Scholes pricing errors after the crash. Of course, the model achieves this by identifying a large and permanent increase in the post-1987 jump intensity – from less than once in 1,000 years to roughly once in 20 years. Because this calibration exercise produces modest pricing errors throughout the sample, a researcher might be tempted to conclude that the market has been efficient over the full sample, and that the jump intensity merely experienced a shift after October 1987. This researcher, of course, would entirely miss the point. The market was not efficient over the full sample. In the initial period, the market price of economic insurance was being offered well below its full cost of production. The jump intensity is merely an *ad hoc* way of modeling this feature of the problem, which allows the researcher to fit market prices, while still failing to notice the actual structure of the situation.

3.3.3 Same Mistakes in a New Setting

We now demonstrate how recent research on CDX tranche pricing focused on seeking small model errors, as opposed to searching for evidence of state-dependence in model errors, allows for a premature conclusion of market efficiency before the learning event, and then after the event, attributes the learning to new information about the state space as opposed to the possibility that the model and those relying upon it were wrong.

A speculative explanation for how this research design is widely relied upon for capital market research is its motivating inspiration from the fundamental theorem of asset pricing – there is no arbitrage if and only if there exists an equivalent martingale measure. This allows for *ad hoc* reduced form Q-measure models to be used to calibrate market prices, and if pricing errors are small (say inside bid-ask spreads), for researchers to conclude that markets are efficient. The notion of state-contingent model errors, their consequences for risk-bearing capital, and their resulting contribution to the full cost of production are unidentified in this setup, potentially in market prices, but potentially not in market prices.

Within this methodology, the failing to notice mistake presents itself as an abnormally low model-implied state price for poor economic states. This methodology is typically applied within a specialized market segment, where we have assumed the relatively sophisticated participants are making the mistake. Finding that pricing errors are small among the securities within this segment confirms that there is a model (or a common set of beliefs) that explains these prices well, but does not establish that the model is correct. This methodology provides weak evidence against the specific alternative that the state prices of poor economic states are significantly higher in other market segments, particularly those that have experienced a learning event.

Unlike the birth of option markets, when credit derivative markets began to take off, there were few empirical studies by academics. In one of the only empirical papers on CDOs to appear in the Journal of Finance (or any of the other top finance journals) prior to the financial crisis, Longstaff and Rajan (2008) fit an arbitrage-free model to investment grade CDX tranches.[7] In their paper, they calibrate parameters governing three processes to explain CDX tranche spreads on each date. They find that the model fits the data exceedingly well – in Table II of their paper, they report RMSEs of under 4 basis points for all but the beginning of their sample. Indeed, their model fits so well that after just two years of data, they are prepared to declare the market "highly efficient."

> "[W]hile the fledging index tranche market may have experienced some inconsistencies in the relative pricing of individual tranches, the market matured rapidly and pricing errors were quickly arbitraged away. [Our] results, however, strongly suggest that the pricing in these markets is highly efficient."

Just as with researchers looking at the index options market 20 years earlier, they find that portfolios identified by popular models as being virtually free of risk are priced as such in the capital market. Calibrations of reduced form Q-measure models indicate that the expected time until a systematic credit event – one that causes 35% of investment grade firms to go bankrupt – is precisely estimated to be once every 763 years under the risk-neutral measure (and therefore far less frequently under the physical measure). Again, entirely ignored is the possibility that portfolios identified as being safe by the popular models being used by market participants will experience painful model errors in bad economic states, and that precisely because it is ignored, this feature of the full cost of production is not yet embedded into market prices.

[7] At the time, Rajan was co-head of fixed-income strategy for Citigroup, the second-largest underwriter of CDOs. He was also David Li's former boss.

Instead of viewing the explosive growth in CDO issuance with a degree of skepticism warranted by the earlier experience with index options and portfolio insurance, it is hailed as an exciting financial innovation that helps "complete the financial market by creating high-credit-quality securities that would otherwise not exist." There is simply no consideration of the possibility that, once again, the explosive growth was somehow related to the fact that sellers of economic insurance were offering prices below their full cost of production.

To draw a further parallel between the learning events of 1987 and 2007, we calibrate a model similar in spirit to the three factor model used by Longstaff and Rajan (2008) to study pre-event CDX pricing, now extending the sample through 2012.[8] To explain CDX prices prior to the learning event, we find the jump intensity of the systematic factor to be very low, consistent with the results reported in Longstaff and Rajan. As can be seen in Figure 13, the model's pricing errors rise during the crisis but then quickly return to their low pre-crisis levels. However, after the learning event, the systematic factor's jump intensity increases dramatically and remains elevated over the entirety of the post-event sample. In particular, the intensity of the systematic jump factor increases from once every 1,000 years on average before the event, to once every 65 years on average after the event. This pattern in the market's updated perception of exposure economic insurance products to poor economic states surrounding the learning event mirrors the pattern around the 1987 event. And again, the interpretation of what type of learning this event represents is strongly influenced by one's prior beliefs about market efficiency.

In a paper written well after the financial crisis, Collin-Dufresne et al. (2012) are able to price CDX tranches with minimal errors during the pre-crisis period using a no-arbitrage model with no tail risk. By calibrating their model to fit equity index option prices and the super senior CDX tranche, they are able to fit average historical spreads across the remaining tranches very well. They then add a systematic jump process to their model. This is not needed to explain pre-crisis spreads, as the pre-crisis risk-neutral likelihood of a catastrophic jump event is estimated to be to be less than once in 1,000 years. From their perspective, the fact that prices contain so little premium for the catastrophic state must mean that this state is

[8] Our model is identical to that of Longstaff and Rajan (2008) but in order to speed convergence we fixed the jump sizes of each of the three processes to their average level reported in Longstaff and Rajan (2008) during the pre-crisis period. Our exercise can be thought of as asking what intensities are required to fit tranche spreads post-crisis when jump sizes are held to their pre-crisis levels.

exceedingly unlikely, "[T]raders in the CDX market are typically thought of as being rather sophisticated. Thus it would be surprising to find them accepting so much risk without fair compensation." Their model continues to explain well the market prices after the crisis, but they find that the jump intensity has now increased by an order of magnitude. This is interpreted by the authors as learning about the state space.

Somewhat remarkably, the authors fail to consider the possibility that this dramatic (and permanent) change in the relative pricing of tranches suggests that pre-crisis prices were wrong. To the extent that "sophisticated traders" of senior tranches were indeed "accepting so much risk without fair compensation" because they failed to notice the possibility of state-contingent model errors, this research design fails to notice this feature of the situation too.

From the perspective of our assumed investor mistake, the facts are interpreted quite differently. The overall pattern of a significant and permanent reallocation of the market's view of how credit risk should be apportioned among the various tranches; the effort to re-evaluate the use of these securities in repo transactions; the issuance of tranched corporate CDOs dropping to and remaining at essentially zero, mirroring the volume patterns of portfolio insurance around 1987; and the emergence of a dramatic increase in the intensity of a systematic jump factor all suggest the learning event was far less benign than investors simply learning about an evolving state space. Instead, the learning should be attributed to the mistaken reliance of relatively sophisticated participants on a model that had *ex ante* evidence of state-dependent model errors. These errors were not noticed because they not explicitly looked for in the methodologies that are relied upon to study these markets. Consequently, these events should prompt a re-evaluation of the testing methods that allowed for extreme confidence to build in models that were discovered to be wrong.

4 Conclusion

In this paper, we present evidence that suggests that financial markets are susceptible to persistent mispricing in new financial insurance contracts. We argue that a specific error – failing to notice the possibility of state-contingent model errors, which will contribute to requirements for risk-bearing capital and to the full cost of production of derivative payoffs – can create a perverse market dynamic that is sustained by market forces akin to a self-fulfilling prophesy (Merton (1948)).

The market for financial insurance offered in the form of short-dated equity index options experienced a dramatic restructuring following the stock market crash of October 1987. The Black-Scholes model based on frictionless capital markets explained well the pricing and capital requirements of index options before this episode, and poorly afterwards. Evidence of state-contingent model errors was available to market participants looking for such evidence, but not so obvious that one was forced to recognize it. The events of October 1987, forced all market participants to notice that their model was wrong, but not necessarily to notice the true nature of the situation. Thus, the self-fulfilling prophesy is broken locally, but allowed to reform in other areas of the capital market.

Twenty years later, the market for financial insurance offered in the form of CDO tranches has undergone a dramatic restructuring in the wake of the 2007-2009 financial crisis, sharing many of the time series patterns as the equity index option market. A newly introduced model of credit portfolio loss distributions was used to identify "safe portfolios" that became highly rated and were priced in capital markets at credit spreads that revealed market participants agreed with the model's assessment. Empirical tests confirmed that financial models explained well the market prices of these securities, and researchers concluded that markets were efficient. Again, evidence of state-contingent model errors was available to market participants looking for such evidence, but not so obvious that one was forced to notice.

The significant state-contingent models errors that were realized during the 2007-2009 financial crisis produced large losses on portfolios previously thought to be safe and forced market participants to notice their model was wrong. In addition to a dramatic relative repricing of CDO tranches, repo funding based on non-government securitizations is restricted, and the issuance of tranched corporate CDOs has essentially gone to zero. The market has clearly learned a lesson. A question that remains is whether this lesson includes an appreciation of the structure of the situation, or if this has again gone unnoticed.

References

Berk, Jonathan B., and Richard C. Green, 2004, “Mutual fund flows and performance in rational markets,” *Journal of Political Economy* 112, 1269–1295.

Black, Fischer, and Myron S. Scholes. 1973. “The Pricing of Options and Corporate Liabilities.” *Journal of Political Economy*, 81(3): 637–54.

Brandt, Michael W, Amit Goyal, Pedro Santa-Clara, and Jonathan R. Stroud, 2005, “A simulation approach to dynamic portfolio choice with an application to learning about return predictability,” *Review of Financial Studies* 18, 831–873

Brav, Alon, and J.B. Heaton, 2002, “Competing theories of financial anomalies, *Review of Financial Studies* 15, 575–606.

Breeden, Douglas T., and Robert H. Litzenberger. 1978. “Prices of State-Contingent Claims Implicit in Option Prices.” *Journal of Business*, 51(4): 621–51.

Coval, Joshua, Jakub W. Jurek, and Erik Stafford, 2009, “Economic Catastrophe Bonds,” *American Economic Review* 99, 628–666.

Collin-Dufresne, Pierre, Robert S. Goldstein, and J. Spencer Martin. 2001. “The Determinants of Credit Spread Changes.” Journal of Finance, 56(6): 2177–2207.

Collin-Dufresne, Pierre, Robert S. Goldstein, and Fan Yang. 2012. “On the Relative Pricing of Long-Maturity Index Options and Collateralized Debt Obligations,” *Journal of Finance*, 67(6): 1983-2014.

Cremers, Martijn, 2002, “Stock return predictability: A Bayesian model selection perspective,” *Review of Financial Studies* 15, 1223–1249.

Gabaix, Xavier, Arvind Krishnamurthy, and Olivier Vigneron, 2007, “Limits of Arbitrage: Theory and Evidence from the Mortgage-Backed Securities Market,” *Journal of Finance*, vol. 62(2), April 2007, p. 557-595.

Garlappi, Lorenzo, Raman Uppal, and Tan Wang, 2007, “Portfolio selection with parameter and model uncertainty: A multi-prior approach,” *Review of Financial Studies* 20, 41–81.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny, 2012, “Neglected Risks, Financial Innovation, and Financial Fragility,” *Journal of Financial Economics*, 104, p. 452-468.

Gervais, Simon, and Terrance Odean, 2001, “Learning to be overconfident,” *Review of Financial Studies* 14, 1–27.

Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartstein, 2014, "Learning Through Noticing: Theory and Experimental Evidence in Farming," *Quarterly Journal of Economics*, forthcoming.

Harrison, J. and David Kreps, 1978, "Speculative behavior in a stock market with heterogeneous expectations," *Quarterly Journal of Economics 92*, 323-336.

Johnson, Timothy C., 2007, "Optimal learning and new technology bubbles," *Journal of Monetary Economics* 54, 2486–2511.

Jones, Christopher, and Jay Shanken, 2005, "Mutual fund performance with learning across funds," *Journal of Financial Economics* 78, 507–552.

Lewellen, Jonathan, and Jay Shanken, 2002, "Learning, asset-pricing tests, and market efficiency," *Journal of Finance* 57, 1113–1145.

Longstaff, Francis and Arvind Rajan, 2008, "An Empirical Analysis of the Pricing of Collateralized Debt Obligations," *Journal of Finance*, Vol. 63(3): 529-563.

MacKenzie, Donald, 2006, *An Engine, Not a Camera: How Financial Models Shape Markets*, MIT Press.

Massa, Massimo, and Andrei Simonov, 2005, "Is learning a dimension of risk?" *Journal of Banking and Finance* 29, 2605–2632.

Merton, Robert C. 1973. "Theory of Rational Option Pricing." *Bell Journal of Economics*, 4(1): 141–83.

Merton, Robert C. 1974. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance*, 29(2): 449–70.

Merton, Robert C. 1998. "Applications of Option-Pricing Theory: Twenty-Five Years Later." *American Economic Review*, 88(3): 323–349.

Merton, Robert K. 1948. "The Self-Fulfilling Prophesy." *The Antioch Review*, 8(2).

Mitchell, Mark L., Todd Pulvino, and Erik Stafford, 2001, "Limited Arbitrage in Equity Markets," *Journal of Finance*, 52(2): 551-584.

Pastor, Lubos, and Robert F. Stambaugh, 2009, "Predictive systems: Living with imperfect predictors," *Journal of Finance*, forthcoming.

Pastor, Lubos, and Pietro Veronesi, 2003, "Stock valuation and learning about profitability," *Journal of Finance* 58, 1749–1789.

Pastor, Lubos, and Pietro Veronesi. 2009, "Learning in Financial Markets," *Annual Review of Financial Economics* 1, 361-381.

Peng, Lin, and Wei Xiong, 2006, “Investor attention, overconfidence, and category learning,” *Journal of Financial Economics* 80, 563-602.

Romer, David, 1993, “Rational Asset-Price Movements Without News,” *The American Economic Review* 83, 1112-1130.

Seru, Amit, Tyler Shumway, and Noah Stoffman, 2010, “Learning by trading,” *Review of Financial studies* 23, 705-739.

Timmermann, Allan G., 1993, “How learning in financial markets generates excess volatility and predictability of stock returns,” *Quarterly Journal of Economics* 108, 1135–1145.

Welch, Ivo, 1992, “Sequential sales, learning, and cascades,” *Journal of Finance* 47, 695–732.

Xia, Yihong, 2001, “Learning about predictability: The effects of parameter uncertainty on dynamic asset allocation,” *Journal of Finance* 56, 205–246.

Table 1. History of Chicago Board Options Exchange Capital Requirements.

Strike	Black-Scholes-Merton Put Price	CBOE Margin Pre 10/1987	CBOE Margin Post 10/1987	CBOE Margin 5/1988	Unlevered Profit (Loss) on October 19, 1987
50.0	0.0	2.0	5.0	10.0	-0.8
52.5	0.0	2.0	5.0	10.0	-1.1
55.0	0.0	2.0	5.0	10.0	-1.5
57.5	0.0	2.0	5.0	10.0	-1.9
60.0	0.0	2.0	5.0	10.0	-2.5
62.5	0.0	2.0	5.0	10.0	-3.1
65.0	0.0	2.0	5.0	10.0	-3.9
67.5	0.0	2.0	5.0	10.0	-4.7
70.0	0.0	2.0	5.0	10.0	-5.6
72.5	0.0	2.0	5.0	10.0	-6.7
75.0	0.0	2.0	5.0	10.0	-7.7
77.5	0.0	2.0	5.0	10.0	-8.7
80.0	0.0	2.0	5.0	10.0	-9.9
82.5	0.0	2.0	5.0	10.0	-10.8
85.0	0.0	2.0	5.0	10.0	-11.6
87.5	0.0	2.0	5.0	10.0	-12.1
90.0	0.1	2.1	5.1	10.1	-12.4
92.5	0.2	2.2	5.2	10.2	-12.3
95.0	0.5	2.5	5.5	10.5	-12.0
97.5	1.2	3.7	8.7	13.7	-11.5
100	2.2	7.2	12.2	17.2	-10.6

Table 2. Regression Analysis of Relation between Back-tested Model Errors and Daily S&P 500 Index Return.

This table reports regressions of daily portfolio hedging errors against the contemporaneous daily realized return of the S&P 500 Index (Return) and the squared return (Return^2). Each day the option supplier holds an option portfolio consisting of one short put option at each moneyness level ranging from 0.5 to 1.0, with an increment of 0.025 (21 options) for each of the closest 3 monthly maturities (a total of 63 options). We calculate hedging errors as: $HE(i, t+1) = modelDelta(i, t) \cdot \Delta SPX(t+1) - \Delta Option(i, t+1)$, where the model delta is calculated using the prevailing "market volatility." We assume that the option supplier contributes 100% of the strike price of each option as collateral against these short positions, such that these positions can be considered unlevered. We estimate the daily market volatility as the 21-day forecast of volatility, based on GARCH(1,1). We estimate the daily Black-Scholes implied volatility to be the GARCH 21-day forecast x 1.18, where this adjustment factor is based on the average ratio of the GARCH 21-day forecast to the VIX (GARCH-to-VIX = 0.85). The daily returns to the hedged option portfolio are calculated as: $R_p(t+1) = \sum_i HE(i, t+1)/Capital(t)$. The mean daily residual for each regression is reported in basis points. The implied-value from the estimated regression evaluated at a Return = -20% is denoted as θ. OLS *t*-statistics are reported in parentheses.

Intercept	Return	Return^2	Adj. R2	Mean Daily Residual (basis points)	θ
		1926.01 – 1982.12			
0.0001	-0.0070		0.001	5.9	-0.2%
(5.11)	(-4.63)				
0.0004	0.0059	-2.55	0.433	5.7	10.3%
(31.51)	(5.22)	(-108.46)			
		1983.01 – 1987.09			
0.0003	0.0041		0.004	2.8	0.1%
(18.46)	(2.30)				
0.0005	0.0084	-3.32	0.741	1.5	13.4%
(58.94)	(9.20)	(-58.38)			
		1988.01 – 2012.06			
0.0001	0.0054		0.002	5.5	0.1%
(6.46)	(3.91)				
0.0004	0.0039	-1.93	0.442	4.1	7.8%
(29.08)	(3.81)	(-69.70)			
		1926.01 – 2012.06			
0.0001	0.0046		0.001	5.8	0.1%
(7.05)	(3.90)				
0.0004	0.0078	-2.64	0.501	5.2	10.5%
(44.30)	(9.25)	(-151.27)			

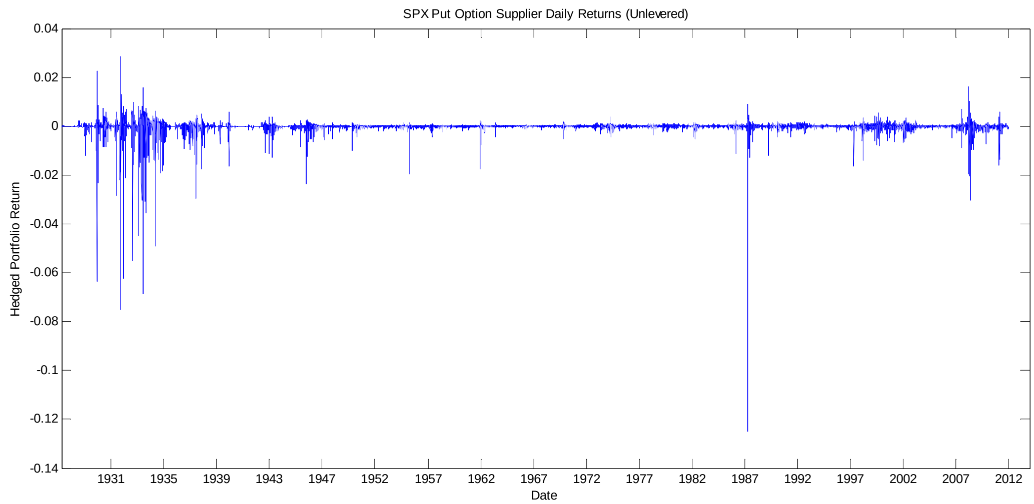

Figure 1. The Experience of an Equity Index Put Option Supplier (January 1927 - June 2012). This figure displays the daily returns to a portfolio of written (short positions) short-dated out-of-the-money index put options that are hedged with the BSM model. Each day the option supplier holds an option portfolio consisting of one short put option at each moneyness level ranging from 0.5 to 1.0, with an increment of 0.025 (21 options) for each of the closest 3 monthly maturities (a total of 63 options). We calculate hedging errors as: $HE(i,t+1) = modelDelta(i,t) \cdot \Delta SPX(t+1) - \Delta Option(i,t+1)$, where the model delta is calculated using the prevailing "market volatility." We assume that the option supplier contributes 100% of the strike price of each option as collateral against these short positions, such that these positions can be considered unlevered. We estimate the daily market volatility as the 21-day forecast of volatility, based on GARCH(1,1). We estimate the daily Black-Scholes implied volatility to be the GARCH 21-day forecast x 1.18, where this adjustment factor is based on the average ratio of the GARCH 21-day forecast to the VIX (GARCH-to-VIX = 0.85). The daily returns to the hedged option portfolio are calculated as: $R_p(t+1) = \sum_i HE(i,t+1)/Capital(t)$.

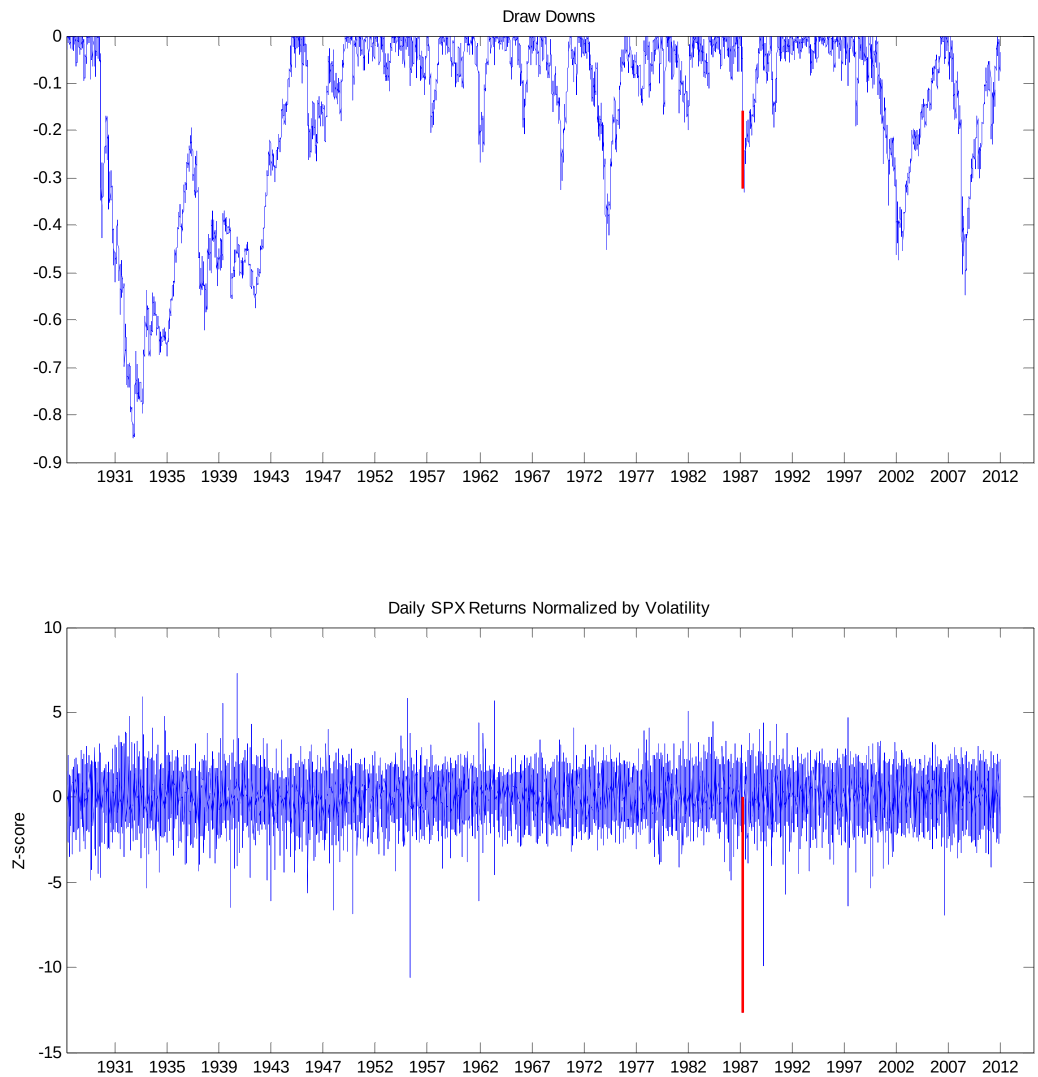

Figure 2. SPX Drawdowns and Daily Z-Scores (January 1927 to June 2012).
Panel A displays the daily level of the SPX relative to its previous historical maximum level as a percentage (drawdown). Panel B displays the daily SPX return scaled by its one-day volatility estimated as of the close of the previous trading day (Z-score). Daily volatility is estimated as GARCH(1,1) using all daily SPX price observations available up to that point in time since January 1926.

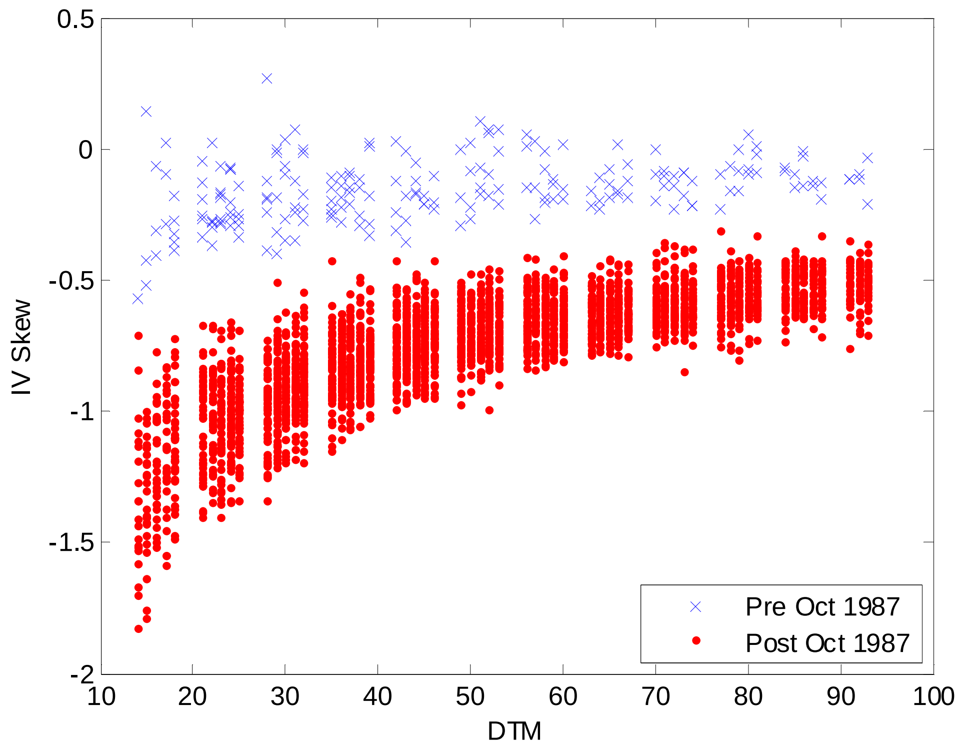

Figure 3a. Short-dated S&P 500 Index Put Option Implied Volatility Skew Pre- and Post-October 1987.
This figure displays the daily slopes of index option implied volatilities as a function of option strike price by the number of days to option maturity (DTM). Each dot represents the slope on a single day. Each day, t, we calculate Black-Scholes/Merton implied volatilities for each put option, IV(t,DTM,K), where the strike price, K, is less than or equal to the prevailing index level, S(t), excluding strike prices more than 2 standard deviations (VIX(t) x sqrt(DTM/365)) below S(t). We estimate the daily IV skew for a given DTM from an OLS regression of IV on K, requiring at least 5 IV observations. Daily IV skew estimates from January 1986 through September 1987 are represented as blue "x"s, while estimates from November 1987 through December 1995 are shown as red "."s.

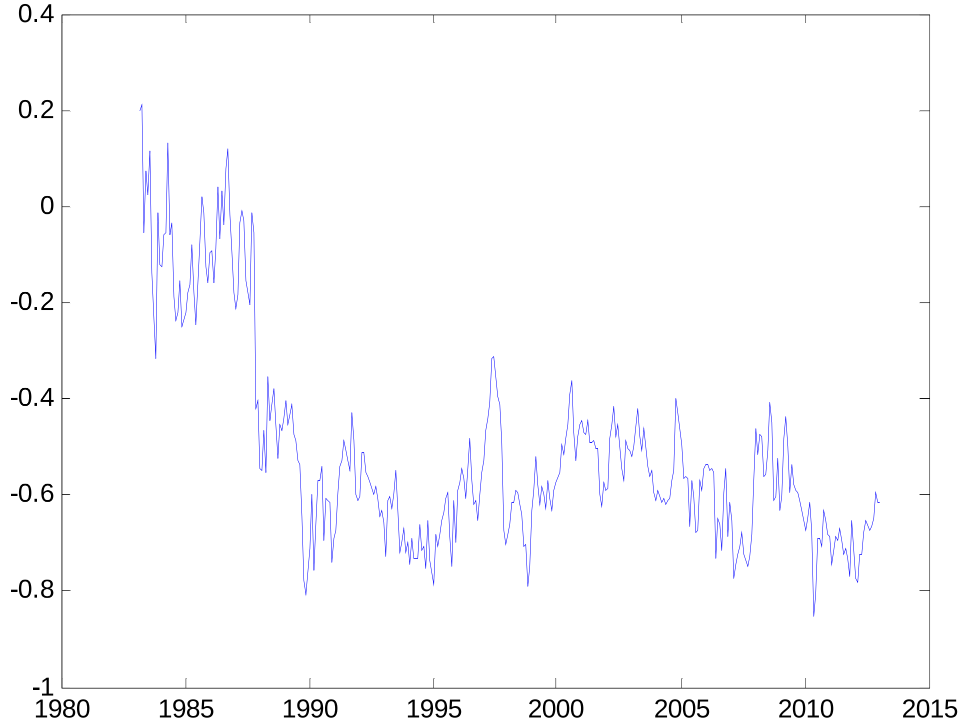

Figure 3b. 60-day Implied Volatility Skew (March 1983 – December 2012). This figure reports the monthly implied volatility skew using all out-of-the-money put options with between 40 and 80 days to maturity. Each day, t, we calculate Black-Scholes/Merton implied volatilities for each put option, IV(t,DTM,K), where the strike price, K, is less than or equal to the prevailing index level, S(t). We estimate the daily IV skew for a given DTM from an OLS regression of IV on K, requiring at least 5 IV observations. The figure reports the median daily IV skew for each month.

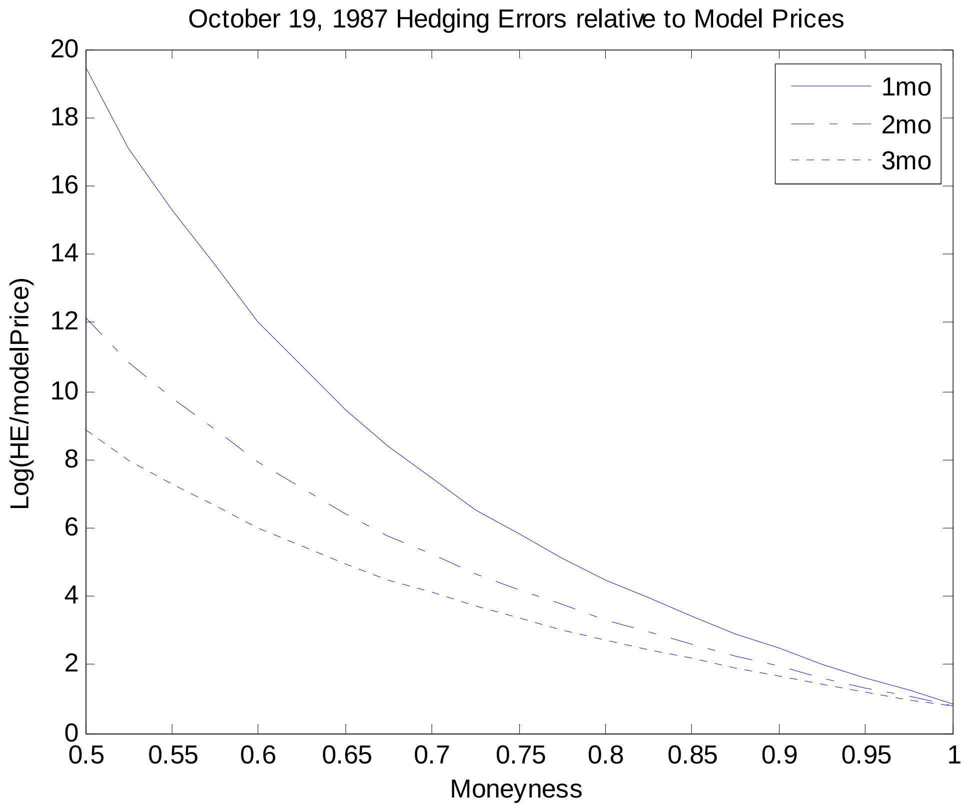

Figure 4. The Realized Experience of an Equity Index Put Option Supplier on October 19, 1987. This figure displays the log of hedging errors as a percentage of the Black-Scholes-Merton model price for each of the holdings in the hypothetical option supplier's portfolio on October 19, 1987. Hedging errors are calculated as: $HE(i,t+1) = modelDelta(i,t) \cdot \Delta SPX(t+1) - \Delta Option(i,t+1)$, where the model delta is calculated using the prevailing VIX as the estimate of volatility.

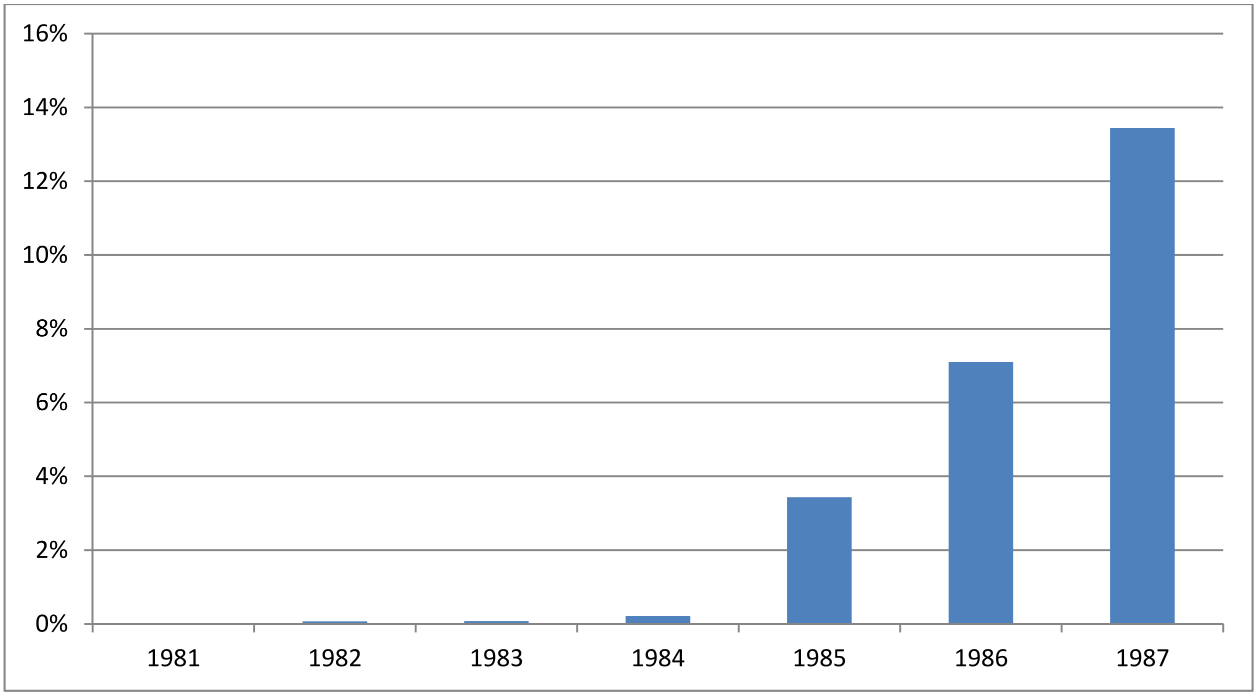

Figure 5. Notional Value of Aggregate Pension Fund Equity Portfolio Protected by Portfolio Insurance. This figure displays data from "Leland O'Brien Rubinstein Associates, Inc.: Portfolio Insurance," *HBS Case Study* by Peter Tufano and Barbara Kyrillos (1995).

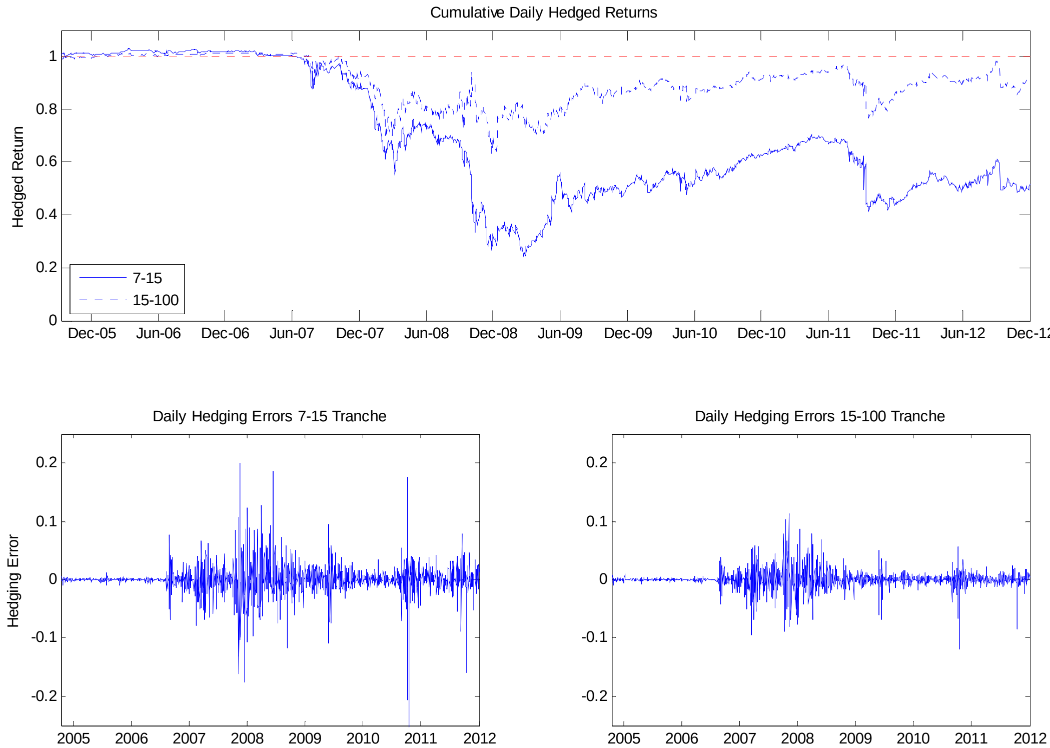

Figure 6. The Realized Experience of a CDO Underwriter (September 2005 through December 2012). These figures report the daily hedging errors for the 7-15 and the 15-100 CDX tranches, hedged with their empirical sensitivity to the underlying CDX. We calculate hedging errors simply as:
$HE(i,t+1) = \widehat{\Delta}(i,t) \cdot chgCDXProtection(t+1) - chgTrancheProtection(t+1)$,
where the empirical delta, $\widehat{\Delta}(i,t)$, is calculated using changes in the all-in upfront prices of protection over the past 100 trading days. To meaningfully compare these dollar hedging errors across securities and through time, we calculate daily hedged returns to the underwriter, assuming that the tranche underwriter contributes capital of $0.20 per $1 of notional tranche exposure for the 7-15 tranche and $0.08 per $1 notional for the 15-100 tranche. The daily returns to the hedged tranche are calculated as: $R_i(t+1) = HE(i,t+1)/Capital(i,t)$.

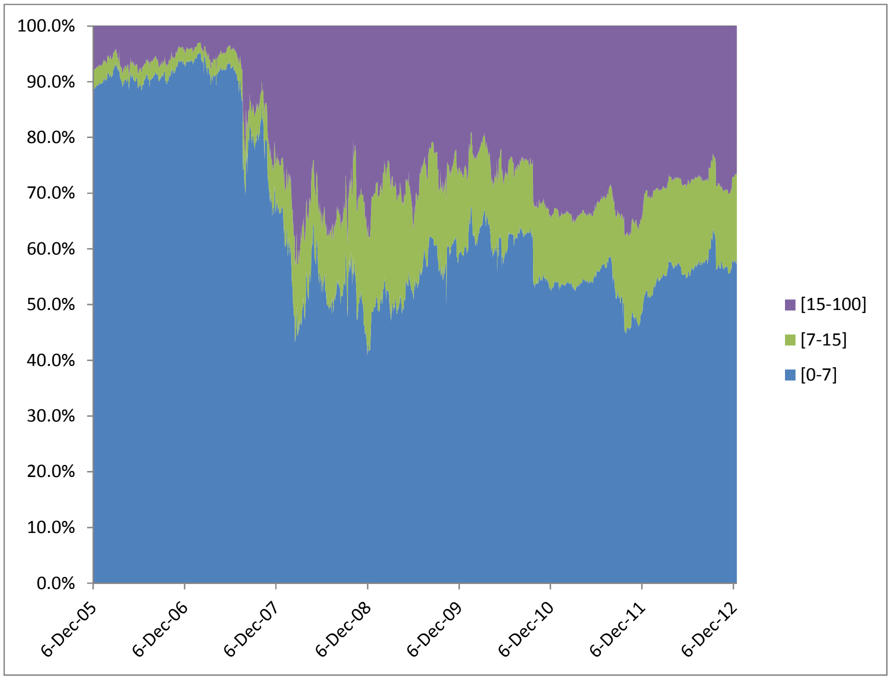

Figure 7. The Market's Allocation of the CDX Risk (September 2005 through December 2012). This figure reports the fraction of CDX yield promised to the Junior (blue), Mezzanine (green), and Senior (purple) CDX tranches.

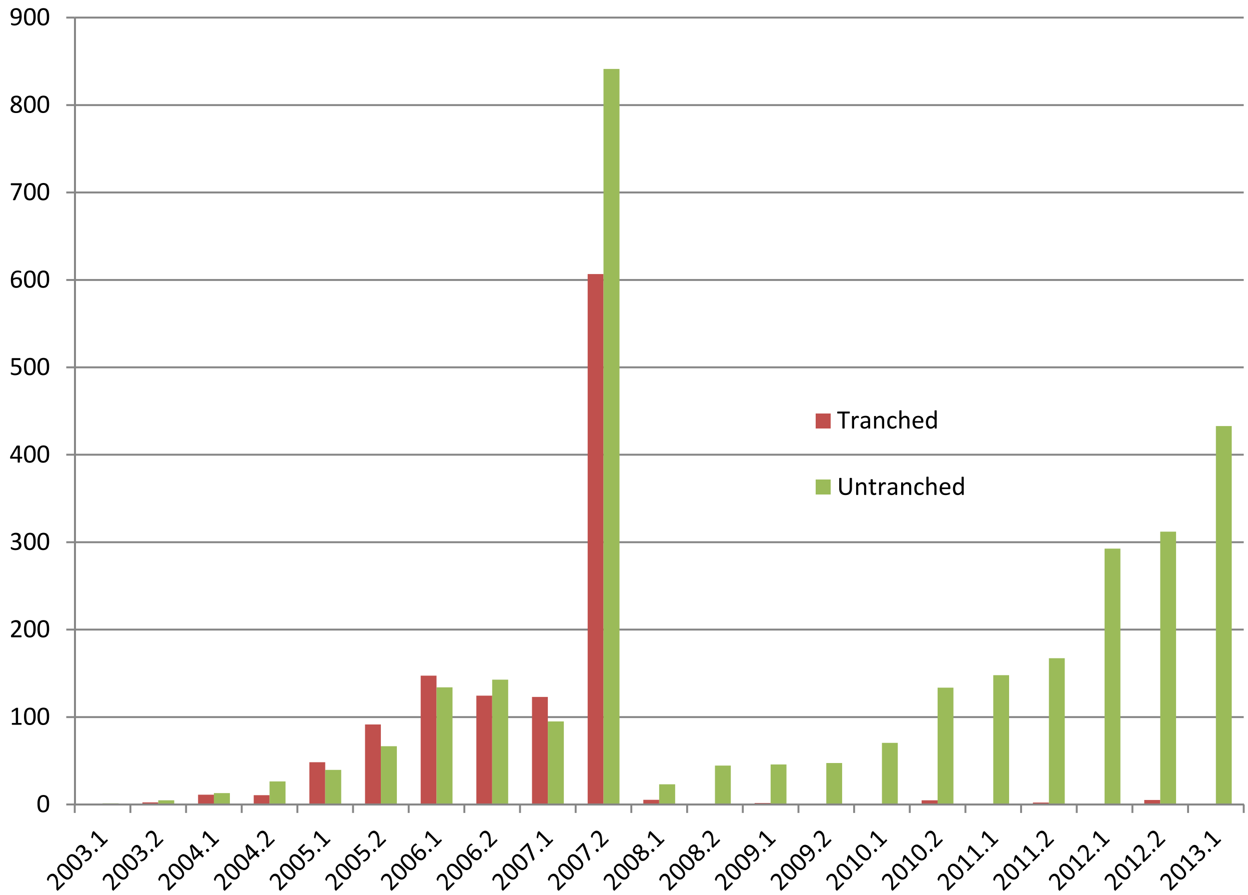

Figure 8. CDX.NA.IG Gross Notional Trading Volume (USD Billions). This figure reports gross notional trading volumes in tranched (red) and untranched (green) CDX.NA.IG contracts according to series issuance date.

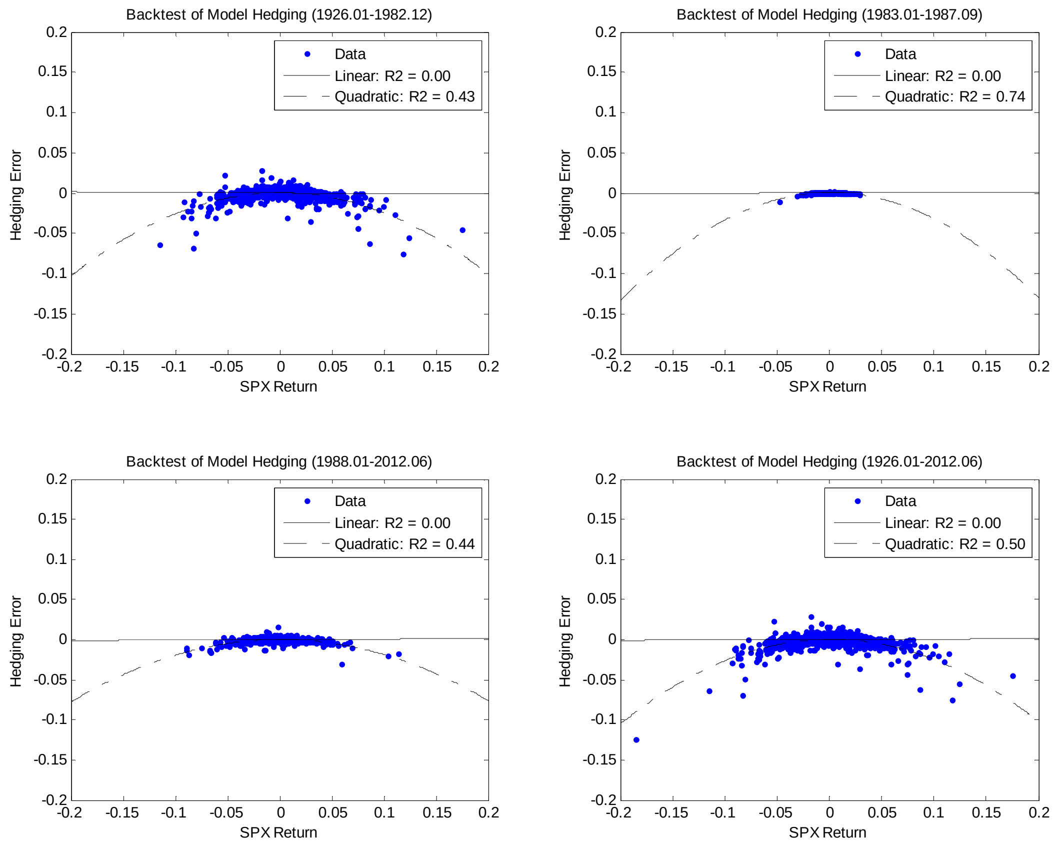

Figure 9. State-Contingent Properties of Model Hedging Errors.

This figure plots daily portfolio hedging errors against the contemporaneous daily realized return of the S&P 500 Index. Each day the option supplier holds an option portfolio consisting of one short put option at each moneyness level ranging from 0.5 to 1.0, with an increment of 0.025 (21 options) for each of the closest 3 monthly maturities (a total of 63 options). We calculate hedging errors as: $HE(i, t+1) = modelDelta(i,t) \cdot \Delta SPX(t+1) - \Delta Option(i, t+1)$, where the model delta is calculated using the prevailing "market volatility." We assume that the option supplier contributes 100% of the strike price of each option as collateral against these short positions, such that these positions can be considered unlevered. We estimate the daily market volatility as the 21-day forecast of volatility, based on GARCH(1,1). We estimate the daily Black-Scholes implied volatility to be the GARCH 21-day forecast x 1.18, where this adjustment factor is based on the average ratio of the GARCH 21-day forecast to the VIX (GARCH-to-VIX = 0.85). The daily returns to the hedged option portfolio are calculated as: $R_p(t+1) = \sum_i HE(i, t+1)/Capital(t)$.

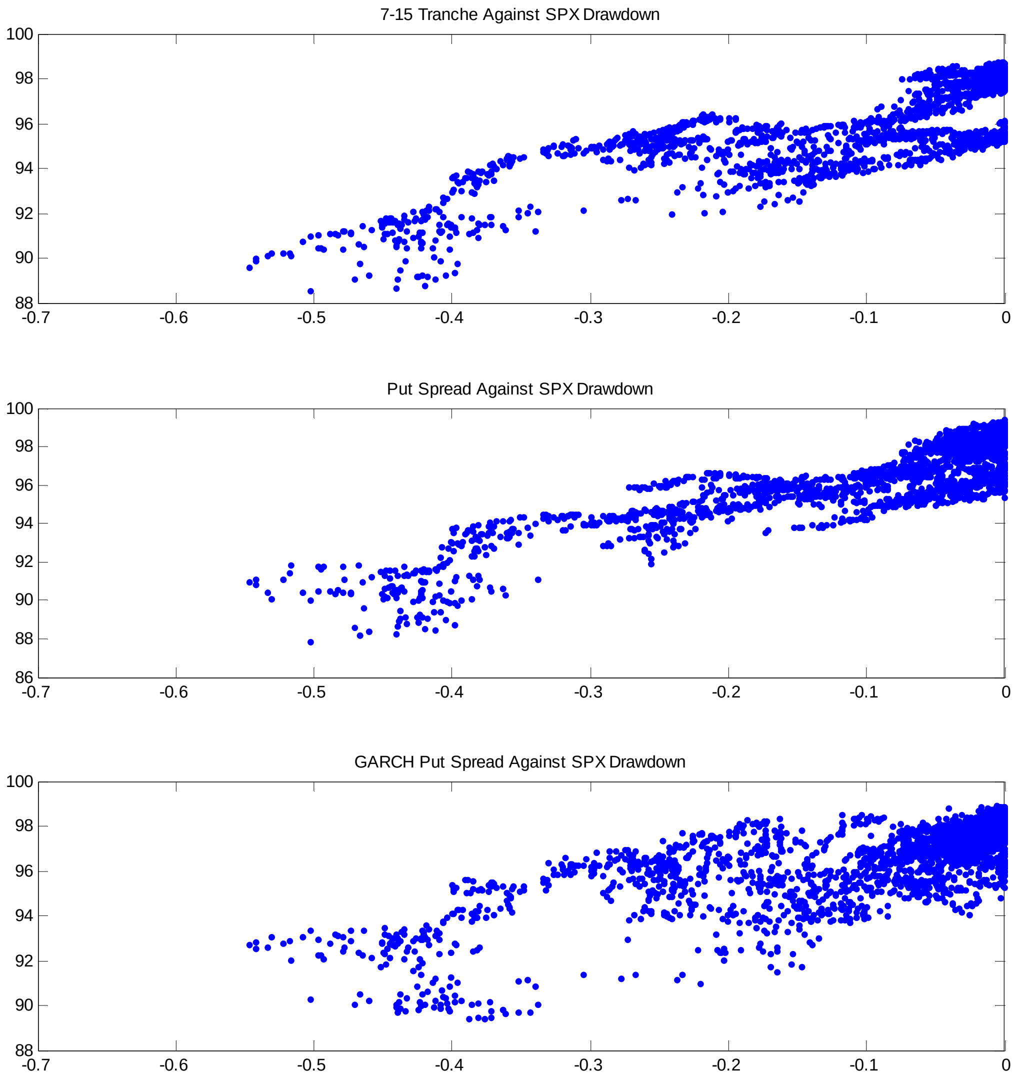

Figure 10. 7-15 CDX Tranche Values across Economic States (2004-2012).

This figure displays the daily values of the 7-15 tranche of the 5-yr North American Investment Grade CDX plotted against the contemporaneous value of the S&P 500 Index (SPX) drawdown. Drawdown is defined as the minimum of zero and the percentage difference between the current index value and its previous maximum value. The top panel displays the actual 7-15 tranche values; the middle panel displays the values of an equity index option replicating portfolio based on actual option prices; and the bottom panel displays the values of an equity index option replicating portfolio based on GARCH-based option prices.

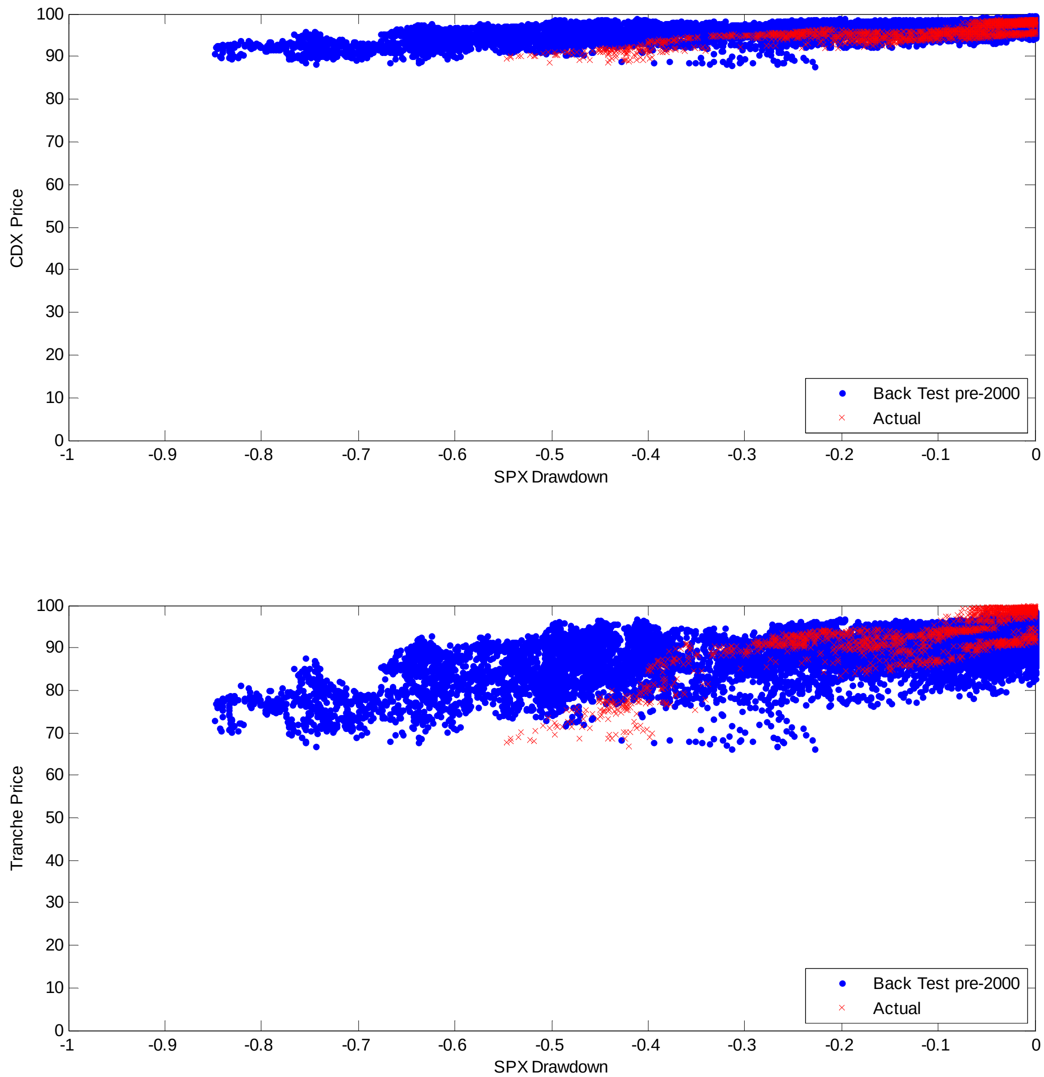

Figure 11. Comparison of Backtested CDX and 7-15 Tranche Values across Economic States (1926-2000).

This figure displays the daily values of the 5-yr North American Investment Grade CDX and the 7-15 Tranche plotted against the contemporaneous value of the S&P 500 Index (SPX) drawdown. Drawdown is defined as the minimum of zero and the percentage difference between the current index value and its previous maximum value. The top panel displays the actual CDX values (red "x"s) from 2004-2012 and backtested values (blue dots) from 1926-2000. The bottom panel displays the actual 7-15 Tranche values (red "x"s) from 2004-2012 and backtested values (blue dots) from 1926-2000. The backtested values are based on equity index option replicating portfolios with option prices estimated from GARCH.

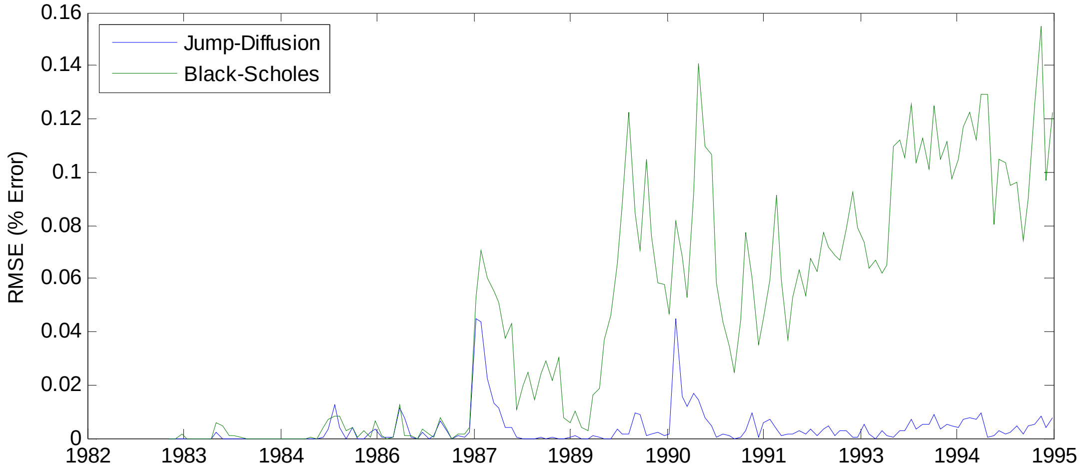

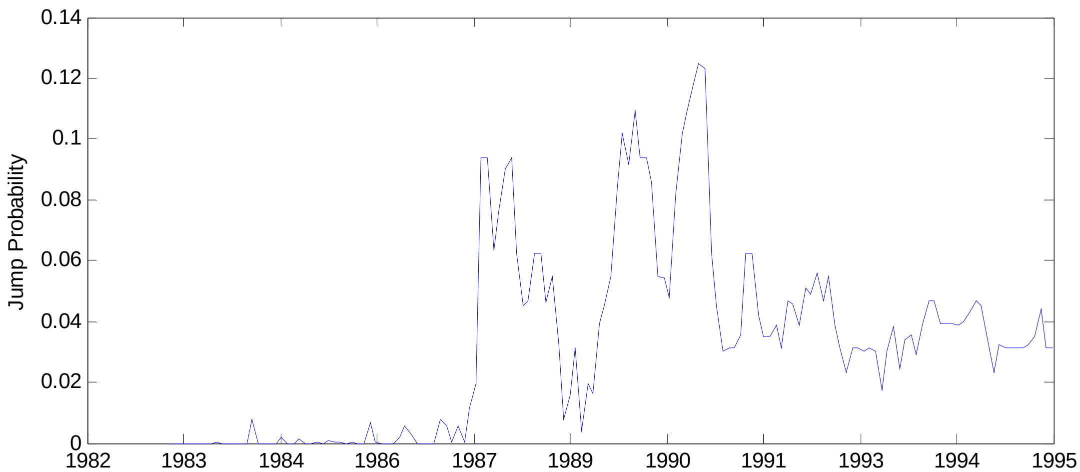

Figure 12. Calibrated Option Model Pricing Errors.
This top figure displays the time series of root mean squared percentage model pricing errors (RMSE). Each day, the Black-Scholes-Merton option pricing model (Black-Scholes) is calibrated to minimize percentage pricing errors between observed market prices and the model price, by choosing the volatility, requiring at least five market prices. Similarly, the Black-Scholes-Merton option pricing model with a jump process (Jump Diffusion) is calibrated to minimize percentage pricing errors between observed market prices and the model price, by choosing the volatility and the jump intensity, assuming a -50% jump size. The second figure plots the calibrated jump intensity. The plotted quantities are the median values calculated from daily observations within a month.

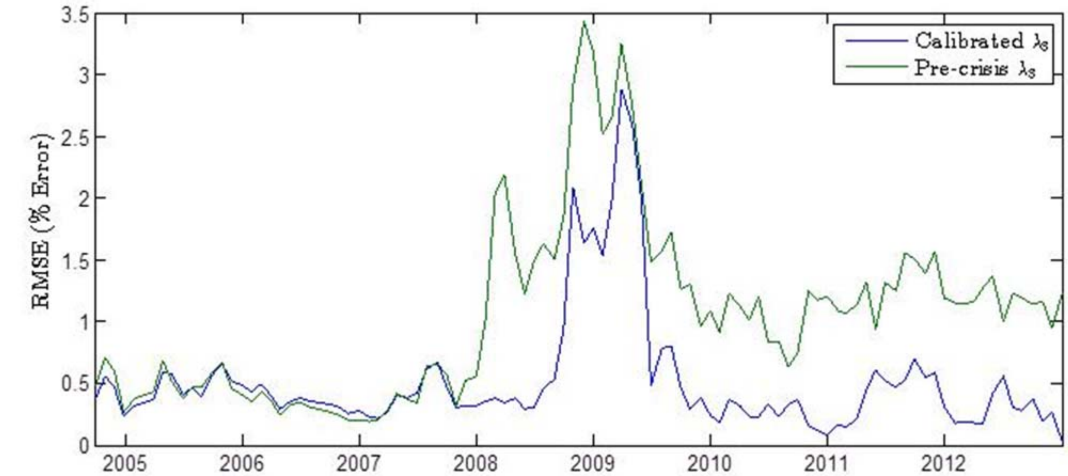

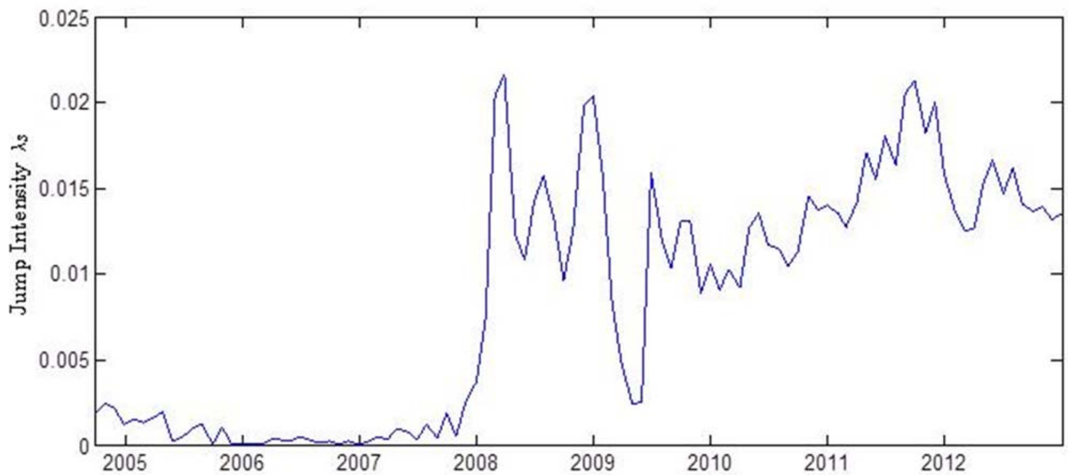

Figure 13. Calibrated CDX Tranche Pricing Errors.
This top figure displays the time series of root mean squared percentage model pricing errors (RMSE). For each CDX series, the Longstaff and Rajan (2008) three-factor portfolio credit model is calibrated to minimize the sum of squared pricing errors (SSE) between the observed market spreads and the model-implied spreads on the equity (0-3%), mezzanine 1 (3-7%) and mezzanine 2 (7-15%) tranches, subject to the model-implied value of the CDX index equaling the market value of the CDX index. Series volatility and weekly jump intensity parameters for each of the three Poisson processes (Calibrated λ_3) are jointly calibrated for each CDX series to minimize the SSE, holding fixed the jump size parameter to the average parameter estimates for CDX 1 – 5 reported in Longstaff and Rajan (2008). Taking these calibrated parameters as given but holding fixed the jump intensity of the third Poisson process to the pre-crisis (09/15/2004 – 08/09/2007) average jump intensity of the third Poisson process yields the second series (Pre-crisis λ_3). The second figure plots the calibrated jump intensity for the third Poisson process. The plotted quantities are the median values calculated from daily observations within a month.