Investment Analysis of Autocallable Contingent Income Securities*

Rui Albuquerque
Boston University, Católica-Lisbon School of Business and Economics
CEPR, ECGI

Raquel Gaspar
CEMAPRE – ISEG, Universidade Lisboa

Allen Michel
Boston University

This Version: December 9, 2014

Abstract

Autocallable contingent income securities, or autocalls, are a relatively new type of structured finance security whose payout is contingent on the performance of an underlying asset and that give investors an opportunity to earn high yields in a low interest environment. We collect data on autocalls issued in the US and describe their contractual properties and the properties of their underlying assets at issuance. We find that autocalls are issued on underlying assets displaying high volatility, negative skewness and high prices. We then model a typical autocall under different assumptions about the price of the underlying asset and (i) analyze the rationale behind the characteristics of the underlying asset at issuance, and (ii) discuss valuation of autocalls in the various models. While the literature consistently finds that structured products are overpriced, we find that incorporating stochastic volatility into the pricing model can help explain some of the overpricing routinely reported in prior studies.

*We thank Kyle Darres for his excellent research assistance, Rachael Christensen and Scott McEwan at Morgan Stanley for providing data and Harvey Boshart, CFA, at JPMorgan Chase for his generous assistance and review of the project. We are particularly grateful to Steve Ross and the Squire Ridge Company, LLC for their financial assistance.

1. Introduction

The low yield environment, driven by four successive rounds of quantitative easing since 2008, has provided the incentive for major financial institutions to develop and sell a new class of structured notes characterized by apparently high yields to the investor. The structured notes market is the fastest growing sector of the US investment-grade fixed-income market (Fabozzi (2005)) and is also fast growing worldwide. Bergstresser (2008) estimates that the total worldwide amount of structured notes outstanding more than doubled every 18 months since 2003 to reach a peak of $4.5 trillion in 2006, dropping to $3.4 trillion in 2008.

This growth not-withstanding, financial advisors that help intermediate the sale of these products often associate the high yield of these securities with their complexity and are careful to promote them outside of their high-net-worth client pool and those sophisticated investors having a view of the market (Purnell (2012)). Recent investment regulation in Europe (European Commission MiFID II Directive), the U.S. (Dodd-Frank Act), the U.K. (Financial Conduct Authority (2013)), and in other countries (see references in Chang et al. (2013)) has pushed for greater investor protection including the separation between independent and non-independent advice, limitations on receipt of commissions and even the prohibition of marketing and distribution of certain products (suitability requirements).

The purpose of this paper is to further our understanding of these seemingly contradictory features of the marketplace for structured products by presenting an analysis of the pricing and return characteristics of these product. There are two main challenges in a study of the pricing of structured notes. First, these are extremely complex products with many payoff features to consider. Second, the panoply of variations on offerings of structured notes implies a pricing model that is almost unique to the issuing security. By virtue of these challenges, important, large-scale studies that speak to the overpricing or underpricing of structured notes (e.g., Bergstresser (2008) and Célérier and Vallée (2014)) cannot distill the causes for the observed pricing from the properties of the notes. In this paper we propose to take a different path, namely to inspect the payoff properties of a specific product called Autocallable Contingent Income Securities, or autocalls. Autocalls have experienced considerable growth in the U.S. market and are a part of a body of structured finance offerings called reverse convertibles, whose payoff is positively linked to the performance of an underlying asset. While our analysis pertains to autocalls, our conclusions have implications for the study of this broader market of structured notes.

Consider as an example the Morgan Stanley July 16, 2012 autocall with Apple, Inc. as the underlying asset.[1] Apple, Inc. stock on July 16, 2012 opened at $605.12, within 7% of its prior all time high of $644 reached on April 10, 2012. With a maturity of three years, the autocall offered a coupon of 3.525% per quarter,[2] (approximately 14% per year) significantly higher than was currently available in the marketplace. The investor receives the coupon only if Apple's stock price on each quarterly determination date is greater than or equal to a threshold level, which in this case is $450 or 75% of the initial share price of $600. In addition, if Apple's stock price on any determination date is greater than the initial stock price, then the security is automatically called and the investor receives the principal. At maturity, if the stock is at or above $450, the investor receives the full amount of the principal and the final interest payment, otherwise, and this is the reverse convertible feature, the investor receives only a fraction of the

[1] In 2012 more than 450 Apple-linked structured products were brought to market, with at least 75% of them issued when the stock was at least $550 (Zweig (2013)). Through mid 2013, we find 68 contingent income autocalls issued on Apple based on SEC securities registration data.

[2] The preliminary prospectus indicates a payment in the range of 3.25% to 4.25% of the stated principal amount per quarter. The final coupon offered was 3.525%. The free writing prospectus can be found at http://www.sec.gov/Archives/edgar/data/895421/000095010312003734/dp31778_fwp-ps257.htm.

principal, namely shares of Apple stock based on a purchase price of $600 or the current cash value of those shares.

We hand collect data on all autocallable contingent income securities listed on SEC's EDGAR from June 2009 to June 2013, totaling $9.6 billion of notional value. In terms of product contractual characteristics, we find that two thirds of all autocalls differ at least in one of the many product features as compared to the plain vanilla example of the Apple autocall.[3] We also find that the median annualized coupon rate is 10% and the median maturity is 1 year. In terms of characteristics of the underlying asset, we find that in two thirds of all autocalls the underlying asset is the stock of a publicly listed company; about 60% of the autocalls are issued at times when the underlying asset's 1-year option-implied volatility (from options available at the date of the issue) is higher than its historical implied volatility after adjusting for movements in volatility of the S&P (a similar result is found in Henderson and Pearson (2011) for SPARQS, interest paying callable notes that are exchanged for shares in the underlying asset upon maturity); about 57.5% of all autocalls have underlying assets that display negatively skewed returns; and, in 50% of the autocalls the price of the underlying security at the issue date is within 12% of the underlying's 52-week high price (a similar result is found also in Henderson and Pearson (2011) for SPARQS).

We value the autocall under three alternative price processes for the underlying asset. We start by assuming that the underlying asset's price follows a geometric Brownian motion,

---

[3] Using a dataset from Bloomberg that includes autocalls issued outside the U.S., Deng, McCann and Mallett (2011) find that the autocall market has reached an annualized rate of more than $30 billion in 2010 and is growing at an average of about 60% per year. Autocalls in the U.K. are frequently based on the FTSE, while those in continental Europe are often based on the Eurostoxx 50 and other major indices from Europe and elsewhere. Like autocalls in the U.S., European autocalls have varied structures, with different and often multiple types of underlying assets, different payoffs, different levels of contingent protection barriers and different degrees of liquidity for the investor. In Japan, Uridash autocalls have been popular instruments for the retail investor and have grown significantly over the last decade partly because they are a cost efficient way for retail investors to engage in the foreign exchange carry trade. These autocalls have benefitted from the low interest rates available to Japanese Yen deposits and the fact that they are denominated in a foreign currency.

which is the workhorse model in the literature. We find that the Apple autocall described above has a fair market price of \$9.86, or 1.4% below the actual price. We then depart from the geometric Brownian motion model in evaluating autocalls—and thus depart from the vast majority of the literature—and study in turn a model of stochastic volatility and a model of mean reversion in prices. We believe that the geometric Brownian motion is not an appropriate model for the underlying assets' price because of several factors. These include (i) our empirical findings, those of Henderson and Pearson (2011), those of Bergstresser (2008), and practitioner discussions (Millers (2013)) regarding the underlying assets' price characteristics at issuance that suggest that underwriters do not choose underlying assets at random; (ii) the large body of evidence of stochastic volatility showing differences in short and long run volatility and, (iii) the vast evidence suggesting reversals in stock prices.

When the underlying asset's price displays stochastic volatility à la Heston (1993), the impact of relatively high volatility at issuance (i.e., short-run volatility) on the valuation of autocalls can differ from long-run volatility, which is not possible under the geometric Brownian motion model where volatility is constant. We show that under reasonable parameters for the Heston model the fair market price of the Apple autocall becomes \$9.98, which implies an overpricing of only 0.2%. Our last model, where the underlying asset's price follows a mean reverting process, allows us to rationalize the systematic use by the underwriters of underlying assets that trade at high prices at issuance. This is not possible under the geometric Brownian motion model where future returns do not depend on the initial price. When prices are allowed to mean revert, we find that the fair market price of the Apple autocall becomes \$9.72, an overpricing of 2.8%.

We conclude from our analysis that the choice of underlying asset or the timing of issuance does not appear to be random and that the pricing of autocalls, and structured products more generally, by adopting models that cannot accommodate changing volatility and mean reversion in prices, can lead to significant biases including overstating the amount of overpricing in these products. These conclusions, however, suggest that there may be clients with certain views regarding price and volatility dynamics on the underlying assets for which these securities are sensible investment vehicles.

Our valuation analysis ignores two important features that unambiguously lead to an increase in the costs to the investor. One feature is the credit risk of the issuer. These securities are backed by the credit of the issuer, not the credit of the underlying asset or that of the distributor. For example, if an autocall is issued by JPMorgan Chase structured with Ford stock as the underlying security, the credit would be that of JPMorgan Chase, not Ford. Because not all banks have the same credit quality and because autocalls are not rated, retail investors are faced with the issue of considering the quality of the issuing bank to determine the value of the security. Pereira da Silva and Silva (2013) in a study of the Portuguese structured retail product market find that this hidden credit cost averages 4.9% per year. Deng, Huali and McCann (2009) find that with the increased borrowing costs faced by Lehman leading up to its distress, the bank issued an increasing number of structured products without compensating the retail investors for the increased credit risk (see also Deng, Mallett and McCann (2011)). The other feature is the potential lack of liquidity of these assets. In the US there is generally little liquidity for the investor prior to maturity or the redemption date, although some issuers provide daily liquidity for the autocalls they have issued. European autocalls also have virtually no secondary market. However, many UK autocalls provide for daily liquidity through marketmaking by the issuer or the London Stock Exchange (London Stock Exchange (2014)). Overall the evidence in this paper supports the regulatory efforts to strengthen investor protection.

We contribute to the literature by studying the properties of the underlying asset's price at the issuance of the structured notes. To our knowledge, only two other papers study these properties. Bergstresser (2008) studies a vast array of notes with call and put like options embedded and Henderson and Pearson (2011), like us, study a specific product, SPARQS.

The literature suggests that traditional (i.e. non-autocallable) structured products are overpriced. Two approaches have been followed. The most popular approach is to assume a model for the underlying asset's price from which the value of the structured product can be derived. The preferred model in the literature is the geometric Brownian motion model (e.g., Burth et al. (2001), Henderson and Pearson (2011), Stoimenov and Wilkens (2005)). To our knowledge only two papers deviate from this benchmark besides ours. Pereira da Silva and Silva (2013) also use a Heston model and Célérier and Vallée (2014) use a local volatility model, but neither paper calibrates their models to reflect the discrepancy between long-run volatility and volatility at the issuance date that we and others find.[4] The other approach values the structured product by replicating its payoff using bonds and options traded (e.g., Burth et al. (2001) and Wilkens et al. (2003)). The advantage of this later approach is that it is model free. The disadvantage is the difficulty in accounting for transaction costs necessary for replication.

Using the geometric Brownian motion as a model for the underlying asset's price, the literature generally finds that underwriters overprice the structured products they sell (e.g., Wilkens et al. (2003), Bergstresser (2008), Henderson and Pearson (2011), Bernard et al. (2011), and Deng et al. (2014)). Whereas we point to an explanation for overpricing that has to

[4] Szymanowska et al. (2009) use a constant elasticity variance model, but as in the geometric Brownian motion model, volatility is constant in this model as well.

do with the initial conditions of the underlying asset and the biases generated by the model used for the price of the underlying, the literature points to clientele explanations and behavioral explanations. Clientele explanations include hedging needs and taxes but also transactions cost explanations that rely on the inability of retail investors to trade in certain markets (e.g., futures and options markets) at the same prices that large institutions can. The general sense is that these are not large enough to explain the findings (e.g., Bergstresser (2008) and Henderson and Person (2011)). Another usual clientele explanation for the demand for high-yielding structured products is the low rate environment and the ability to achieve some degree of capital protection (e.g., Burth et al. (2001), Stoimenov and Wilkens (2005), Coval et al. (2009), Szymanowska et al. (2009), and Stein (2013)).

Several behavioral explanations have been advanced including investor irrationality or bounded rationality, framing, and overweighting of small probability events (see Breuer and Perst (2007), Bergstresser (2008), Hens and Rieger (2008), Bernard et al. (2011), Henderson and Pearson (2011), Das et al. (2013)). Interestingly, studies of structured finance securities trading in the secondary market reveal that the overpricing disappears over time and pricing reverts to the theoretical price of the security (see Stoimenov and Wilkens (2005)). Product complexity is often advanced as a behavioral explanation to explain the cross-sectional variation in prices in this market (Stoimenov and Wilkens (2005), Pereira da Silva and Silva (2013), and Célérier and Vallée (2014)), but it is less clear why it would predict overpricing.

Section 2 presents a description of the properties of autocalls and Section 3 presents the model of the autocall that we study. Sections 4, 5 and 6 present the autocall pricing results when the underlying asset's price follows a geometric Brownian motion, the underlying asset's price displays stochastic volatility, and the underlying asset's price is mean reverting, respectively. Section 7 concludes. The Appendix contains the mathematical formulations of the models.

2. Autocall Sample Characteristics

2.1 Characteristics unrelated to the underlying asset

We collect the universe of contingent income autocallable securities from the SEC's EDGAR database from June 18, 2009 through June 4, 2013. We search all the prospectuses during the observation period using form 424(b)(2) and the search terms "autocallable" and "contingent income" resulting in 1,162 autocalls. Table 1 provides summary statistics of our autocall data. Panel A shows that the number of autocalls has increased significantly over time with the principal value over the period exceeding nine billion dollars.[5] Panel B of Table 1 lists the underwriters. There are two main bank underwriters, JPMorgan Chase and Morgan Stanley, representing 73% of all issuances and 57% of the issued dollar volume.

[INSERT TABLE 1 HERE]

There are many variations on the plain vanilla autocall described in the Introduction. The standard, or plain vanilla autocall has a single underlying asset, has a fixed coupon rate, can be automatically called if certain conditions hold, has a fixed threshold below which the coupon is not received, has a fixed final payment conditioned on whether the underlying asset is above

[5] The SEC search only permits a four-year historical window to be searched. Prior to 2009, using Bloomberg, we were able to find 29 autocalls registered with the SEC representing $226 million. Because prior to 2009 there are so few autocalls, we constrained our data to the autocalls having form 424(b)(2) available on EDGAR to assure ourselves of the comparability and quality of the data provided. Deng, Mallet and McCann (2011) instead search the Bloomberg database for autocalls and are able to obtain a significantly larger sample than ours. Their larger sample is explained mostly by the inclusion of autocallable securities issued outside the U.S.

or below the threshold level, and has the possibility of early redemption. Panel C of Table 1 shows that roughly 2/3 of all autocalls differ in at least one feature from the vanilla autocall.

Panel D of Table 1 shows that 25% of all autocalls in the sample, or 38% of the total principal value, have a threshold of 80%. Panel E presents the coupon rates, expressed on a per annum basis, by percentile and their distribution moments. As the autocalls were developed to provide investors with an opportunity to generate above market yields, it is not unexpected that the median coupon is as high as 10% p.a. and that even at the 25th percentile, the autocall coupon, 8.50%, is still significantly higher than corporate bond yields. The positive skewness of 1.44 indicates that the coupon distribution has a fat positive tail. Figure 1 plots the frequency distribution of the coupon rates showing that the right tail extends to a coupon rate of 33% p.a. Panel F of Table 1 shows that the mode and median maturity is one year, with 43.3% of the autocalls having a one-year maturity. The next most frequent maturity range is between 1.01 and 4.99 years, with 23.2% of the autocalls in the sample.

[INSERT FIGURE 1 HERE]

2.2 Characteristics related to the underlying asset

Panel G of Table 1 shows that the most common underlying asset is an individual company stock (e.g. Apple, Inc.), representing 64.7% of all issues, or 75% of the principal amount underwritten. The next most frequently underwritten securities are equity indices and commodities. Figure 2 identifies the most frequent underlying assets to be the Russell 2000 Index (17.6%), the S&P 500 Index (15.4%), and Apple, Inc. (5.9%).

Panel H presents statistics computed based on the stock's volatility relative to its historical volatility. We calculate the ratio (underlying asset's implied volatility / underlying asset's historical implied volatility) / (S&P's implied volatility / S&P's historical implied volatility). This ratio is used to investigate whether the securities that are chosen for autocalls have high volatility, controlling for changes in aggregate volatility. We find that in about 60% of the cases volatility is higher than historical, with 25% of the autocalls having ratios greater than 1.08 and 25% having ratios less than 0.97. This result can also be observed in Figure 3, which provides the complete frequency distribution of the above described volatility ratio.[6] Our result complements the evidence in Henderson and Pearson (2011) for SPARQS that volatility (weakly) affects positively the choice of the underlying security, but differs from the evidence in Bergstresser (2008) that uses a sample of all structured products and finds mixed evidence on the effect of volatility on the likelihood of the choice of the underlying security. Our evidence thus suggests that the underlying securities are generally very volatile at the time of issuance. We cannot identify the cause for the higher volatility, whether firm specific or market-wide. Indeed, practitioners often suggest that these securities are only issued at times of high market volatility (e.g. Millers (2013)). We use this evidence below to justify the consideration of models with stochastic volatility to price autocalls.

[INSERT FIGURES 2 AND 3 HERE]

Panel I of Table 1 presents the percentiles and distribution moments associated with the ratio of the price of the underlying at the time of issuance of the autocall to its prior 52-week high. This analysis is performed for those securities with a single underlying asset (e.g. the Russell 2000 or Apple stock). The analysis helps gauge how closely the underlying securities' prices are to their most recent highs. A ratio of 1 reflects the fact that an offering was issued precisely at the 52-week-high price. We find that 25% of the autocalls are issued within 4% of

[6] The general tone of these results is unchanged if we use realized historical return data and robust measures of variance, the interquartile range and the median absolute deviation. Results are available upon request.

the 52-week-high price and 50% are issued within 12% of the 52-week-high price of the underlying security. Figure 4 graphically depicts the frequency distribution of the ratio of the price of the underlying security at the time of issuance to its prior 52-week high. This evidence demonstrates that underwriters appear to choose to issue autocalls on underlying assets whose prices at the issue dates are near their 52-week high levels. We repeat the analysis using the variable (initial price of underlying security / underlying security's 52 week high) / (S&P 500 value at issuance / S&P 500 52-week high) that adjusts for market movements. The histogram of this variable using our autocall data looks very similar to that obtained in Figure 4, suggesting that markets have been rising during the time of issuance.[7] The fact that markets have been rising during the time of the study supports the view that there are an increasing number of autocalls issued during periods of market strength.[8] Similarly, Henderson and Pearson (2011) find an increased likelihood of choice of an underlying asset with high past 12-month performance. This evidence leads us to consider models of mean reversion in prices when pricing autocalls. Bergstresser (2008) finds mixed evidence regarding the past performance of underlying assets at issuance. We present a possible explanation for this discrepancy in Section 6. Finally, Panel J shows that mean skewness is negative at -0.31 and 57.5% of all underlying assets display negative skewness. This is striking as firm level stock returns are vastly positively skewed (see Albuquerque (2012)).

[INSERT FIGURE 4 HERE]

3. Framework for Analysis

[7] Available from the authors upon request.

[8] Oyo (2013) suggests that the improved prospects of the U.S. economy following the crisis drove usage of U.S. indices as underlying assets, where the deteriorating prospects of the European economies drove usage away from the more stagnant European-based indices (see also Thin (2014) on market shares).

The contingent income autocallable security that we analyze is of the plain vanilla kind. Payouts are a function of the price performance of an underlying asset as described next (the Appendix gives a formal description of the payouts). At each determination date the security is called if the underlying asset's price at that date is higher than the price at issuance, in which case the investor gets paid the coupon and principal and no further cash flows, or is not called, in which case the investor gets the coupon if the underlying asset's price at that date is higher than the threshold level or zero otherwise. The threshold level is defined as a fraction, say 80%, of the price at issuance. At maturity, if the autocall is still alive, it pays coupon plus principal if the underlying asset's price at that date is higher than the threshold, otherwise it pays either one unit of the underlying asset or its current cash value. If the underlying asset is received, the investor has a capital loss.

Without the autocall feature, this structured product is best described as a combination of a long position in a plain vanilla bond with fixed coupon payments at every determination date and redemption of the principal at par at maturity, plus several short positions. These short positions include several European digital options, each maturing at a different determination date, and one European digital option and one European put option both with a maturity that coincides with the maturity of the autocall. Because of the value associated with all the embedded options given to the underwriter by the investor, which expose the investor to the downside risk of the underlying asset but not to its upside potential, the underwriter is able to offer a higher coupon rate.

The autocall feature significantly complicates this structure because it makes the value of the embedded options contingent on the price of the underlying stock at each of the determination dates. However, it is still true that the investor gives contingent options to the underwriter and the underwriter can use the value of these options to offer a better coupon rate to the investor.

4. Geometric Brownian Motion Model of Underlying Asset Prices

In this section we present results describing the properties of the autocall under the assumption that the price of the underlying asset follows a geometric Brownian motion. Under this assumption, and without the autocall feature, it is possible to write the value of the security in closed form. However, in the presence of the autocall feature there is no known exact, closed-form expression for the value of the autocall security and we must, therefore, resort to numerical methods to determine the properties of the investment return. For concreteness, we simulate the properties of the plain vanilla autocall using the Apple autocall described in the Introduction. We calibrate several model parameters: the maturity of the contract is set to 3 years; the contract has quarterly determination dates and pays an effective coupon rate of 3.525% quarterly; the threshold level in the contract is 75% of the initial stock price; the risk free rate is set to 1.8% p.a., which equals the average annualized nominal 3-month T-bill rate in the period 2005-2012; the mean growth of the price process is set to 6.3% p.a., which can be decomposed as the sum of a risk premium of 4.5% and a risk free rate of 1.8%; and the volatility of the price process is set to 30%, which is the implied volatility of Apple's stock return on options with 1 year maturity at the time of the issue.[9] The investors' required rate of return on the autocall is set to 6.12%. This rate equals the risk free rate of 1.8% p.a. plus the risk premium of 4.5% p.a. times Apple's CAPM beta, that at the time of the issue was equal to 0.96 relative to the S&P 500. The investors' required rate of return is only used to compute the fair

[9] To compute implied volatility we chose 1 year for the option maturity instead of 3 years to match the maturity of the autocall due to a lack of liquidity on long-dated options.

market price of the autocall and as a benchmark for the internal rate of return of the autocall. For our purposes the initial stock price is arbitrary and is set to 10. We simulate 50,000 price paths to ensure the accuracy of our results. With these parameters the fair market value of the security according to our model should have been $9.86, representing an overpricing of 1.4% relative to the actual sale price. The corresponding unconditional expected annualized internal rate of return of this autocall is 4.3% p.a., below the assumed required rate of return of 6.12%.

4.1 Survival probabilities

Table 2 presents the simulated unconditional and conditional probabilities of survival. The row indicated in bold "Baseline case" refers to simulations that use the parameters defined above. To understand the relevance of the various parameters to the value of the autocall, the table also gives the unconditional and conditional probabilities of survival under several other models: low and high volatility (respectively, 15% and 40%), low and high threshold level of the initial stock price (respectively, 60% and 85%), short and long maturity (respectively, 1 year and 15 years), and low and high coupon (respectively, 8% p.a. and 25% p.a.). These alternative models were constructed based on the evidence presented in Section 2.

Consider the unconditional probability that the security is called at the first determination date, which occurs if the stock price at that determination date is above the initial stock price. In the baseline case seen in Table 2, this probability is 51.2%. The probability that the security is called at either the first or second determination dates is 64% (equal to 51.2% plus 12.8%). To understand these numbers, note that with our calibration the mean of arithmetic returns is equal to 0.018. Since the mean of arithmetic returns is positive, the probability that the stock price is larger than the initial stock price at the first determination date is greater than

50%. Therefore, investors with a view that the price will grow will have an incentive to invest in these securities given the high yield paid. The probability that the security is called at the second determination date equals 12.8% (equal to the probability that it is not called by the first determination date, 1-0.512, times the conditional probability that it is called by the second date, given that it was not called at the first date, 26.4%). The conditional probability of being called by the second determination date is considerably lower than 50% because a condition for surviving the first determination date is that the price be strictly lower than the initial stock price. Further, we note that the unconditional probability that it is called at any determination date after the first year is quite low at 3% or less. The probability of not being called at any determination date prior to maturity, namely the probability of reaching maturity, is 15% in the baseline model calibration.

Consider now the scenario where stock price volatility is higher and equal to 40% and all else remains equal. The increase in volatility lowers the mean of arithmetic returns to negative 0.017, which lowers the probability that the security is called at the first determination date to 49.2%. In general higher stock price volatility lowers the probability that the security is called at any determination date and raises the probability that the security reaches maturity. The fact that the security is more likely to reach maturity also implies that the stock price is more likely to be below the threshold level at maturity and hence that the investor takes a capital loss.

The remaining scenarios considered in Table 2 consider different threshold values, different maturities and different coupons, but share with the baseline case the same probability that the security is called. This is because these dimensions of the autocall do not change the price process or the price level that triggers the autocall. They will, however, impact the value of the autocall as explained next.

[INSERT TABLE 2 HERE]

4.2 Internal Rate of Return

To further analyze the return properties of the autocall we calculate *ex-post* internal rates of return. That is, for each simulated price path, we calculate the corresponding IRR. These are *ex-post* IRRs because they are calculated based on specific realizations of the stock price. These rates naturally differ from the *ex-ante* IRR reported above that was calculated using expected cash flows across all simulated paths. In addition, the average of the *ex-post* IRRs is also different from the *ex-ante* IRR because the IRR results from a nonlinear calculation.

Because of the autocall feature a security that is called at the first determination date has a maturity one quarter shorter relative to another security that is called at the second determination date, and so on. To deal with the issue of heterogeneity in the effective maturity of the cash flows across simulated paths we assume that after the security is called the notional value from the autocall is reinvested at the risk free rate through the maturity of the autocall.[10]

Figure 5 displays properties of the IRR associated with the Apple autocall (baseline case) and with the other scenarios described above. The top row (panel A) in Figures 5A to 5D is common as it displays the results for the baseline case. Figure 5A also depicts the low and high volatility scenarios, Figure 5B depicts the low and high threshold level scenarios, Figure

[10] Alternatively, we calculate the ex-post IRR without assuming reinvestment of capital after the security is called. Quantitatively, IRRs calculated in this fashion are higher than those reported in the paper because the risk-free rate is lower than the coupon rate used to calibrate the model, but qualitatively the properties of both IRRs are quite similar. Because the reported measure is more relevant for investors and to conserve on space we omit the presentation of the results under this alternative approach to calculate the ex-post IRR.

5C depicts the short and long maturity scenarios, and Figure 5D depicts the low and high coupon scenarios.

The first column in each figure presents the unconditional frequency distribution (across all simulated paths) or histogram of IRRs. In each of the histograms in all panels of Figure 5A-5D the tallest bar is explained by the probability that the security is called at the first determination date, pays the coupon and then earns the risk-free rate up to maturity. To its right, the histogram depicts the events where the security was not called until at least the second determination date and is likely to have paid coupons until called.

The riskiness of the autocall can be seen in the significant left tail of the histogram. The distribution of ex-post IRRs is considerably left skewed with non-negligible probabilities of extremely low ex-post IRRs. In the baseline case the IRR is negative in 10.6% of the simulated price paths and is below -5% in 9.6% of the simulated price paths. The negative skewness is more pronounced when volatility is high (negative IRR in 15% of paths and IRR below -5% in 14.4% of paths), the threshold level is high (negative IRR in 12.5% of paths and IRR below -5% in 11.4% of paths), the maturity is shorter (negative IRR in 11.8% of paths and IRR below -5% in 11.8% of paths), or the coupon is lower (negative IRR in 10.9% of paths and IRR below -5% in 10.6% of paths).

[INSERT FIGURE 5 HERE]

To better understand the negative skewness of the IRR distribution consider column two in Figures 5A-5D. The plots in this column present the average IRR across sample paths conditioned on the security surviving a given determination date. The left-most bars in the plots, indicated by determination date 0, give the unconditional mean of the ex-post IRR across all paths. In the baseline case the unconditional mean IRR is 1.91% annually.

There is a generalized downward pattern of the mean IRR when measured against the determination date survived by the security. Intuitively, a necessary condition for the security to have survived each past determination date is for the stock price to be below the initial stock price at each of the determination dates. In the geometric Brownian motion model, because the mean return is constant, the expected value of future prices decreases as the price goes down. Therefore, the expected payoff to the investor at maturity at that time is also expected to be lower. The mean IRR of autocalls that reach maturity is always low because the investor bears the downside of the stock price at maturity. This pattern of IRRs explains the negative skewness in the distribution of IRRs and the reported overpricing. As a result, the investor incurs the risk of significant losses should the autocall not be called soon after issuance, with the mean IRR approximately -10% per year for those autocalls not called by the 11th determination date (33 months after issuance).

If the stock price displays higher volatility, keeping all else constant, the probability that the security is called at the first determination date decreases (see Table 2). Therefore the likelihood that the investor is paid the coupon decreases. Also, from Table 2, the probability that the security survives till maturity increases. Because higher volatility in stock prices increases the probability of prices being below the threshold level, the likelihood that the investor takes a capital loss at maturity increases. IRRs decrease relative to the baseline case.[11]

[11] In the low volatility case, the mean IRRs display a hump-shaped pattern against the determination date. The reason is that with low volatility, the stock price remains closer to the initial price and if the security is not called, then the investor receives the coupon, which is better than having the security be called earlier on and having to reinvest the proceeds at the risk-free rate as assumed. Eventually, if the security is not called, the probability that the price is below the threshold is large enough putting downward pressure on the IRRs and giving rise to the decreasing part of the hump.

The effect of the increased likelihood of reaching maturity and of the price of the underlying being below the threshold level significantly reduces the mean IRR to approximately -20% per year at the 11th determination date.

Figure 5B depicts the effects of changing the threshold level. Consider panel C first. The high threshold level does not affect the probability of the security being called (see Table 2), but strictly lowers the cash flows from the autocall. This is because there are price paths that would generate payouts at certain determination dates that no longer generate those payouts. This results because the threshold level is now higher than the price at those determination dates. The IRRs therefore uniformly decrease relative to the baseline case (panel A), all else constant. In contrast, in the low threshold case (panel B) the fact that more price paths now involve a payment of a coupon partly offsets the negative effect of the capital loss at maturity and generates an increasing pattern in IRRs for the first few determination dates.

Figure 5C depicts the effects of changing maturity. Consider first panel B. The shorter maturity does not affect the probability that the security is called but increases the likelihood of it reaching maturity (see Table 2). While the probability that the stock price is below the threshold value at a certain date is not affected by the maturity of the asset, the increased likelihood of reaching maturity increases investors' expected capital loss and lowers IRRs. This effect is quantitatively very large with the IRR at maturity reaching a whopping -6% per year. This compares with an IRR in the baseline case of about -3% per year at the same date (see panel A). Extending the maturity of the autocall to 15 years (panel C) creates the possibility of very high IRRs at the early determination dates because the effect of the capital loss at maturity is quite far off. But again, a large negative IRR arises as the security reaches maturity. Finally, lowering the coupon (Figure 5D) also strictly lowers the payout relative to the baseline case and shifts all mean IRRs down.

The value of the capital loss faced by the investor at maturity can be further assessed. Consider the worst possible outcome for the investor that buys an autocall: it results from the stock price being below the threshold level at all determination dates and also at maturity in which case the investor receives the underlying stock and earns a capital loss. The return in this case can be easily computed as the IRR of the following present value:

$$-10 + \frac{S_T}{(1+IRR)^{12}} = 0,$$

where S_T is the stock price at maturity T. The solution is approximately equal to $IRR = \frac{1}{12}\log(S_T/10)$. The expectation of this value is difficult to calculate analytically because it requires knowledge of a distribution that is conditioned on the price being below the initial price at every determination date. However, the fact that the stock price at maturity can be close to zero and the fact that the probability of receiving the stock at maturity is high (10.9% in the baseline case) help explain the negative skewness in the return distribution.[12]

In summary, we find that the distribution of ex-post IRRs is highly left skewed; that left skewness reflects the investors' capital loss if the security reaches maturity that is embedded in the short option positions; and, that left skewness is greatly affected by several model parameters including the volatility of the underlying asset. Investors therefore have much to gain in the autocall that is called early in its life and to lose in the autocall that survives to maturity.

[12] The best possible outcome for the investor is to have the stock price lie between the threshold level and the initial stock price at every determination date and for the stock price to be above the threshold level at maturity. In this case, the investor receives the coupon every determination date and principal plus coupon at maturity, equivalent to a 3.525% quarterly IRR over three years in the baseline case. This event is unlikely, occurring with a probability of 0.2% in the baseline model.

5. Volatility Exposure through Autocalls

An important assumption of the model in Section 4 is that of constant volatility. This assumption simplifies the problem at hand but is at odds with the volume of evidence on stochastic volatility in asset returns. In the presence of stochastic volatility autocalls become vehicles to obtain exposure to volatility risk, and pricing of an autocall must differentiate between short-dated volatility and long-dated volatility because of the combination of options of different maturities that are embedded in the autocall. To preview our results in this section, issuance of autocalls at times of high volatility on the underlying asset tends to produce higher valued autocalls if this volatility is expected to decrease over time.

To operationalize these ideas we simulate the stochastic volatility model developed by Heston (1993) (the Appendix contains the model details). This model distinguishes between the conditional variance of stock returns, which we can label as short-dated variance, and the long-term mean of the conditional variance of returns, which we can label as long-dated variance. The process for variance has mean reversion and shocks to variance, with the shocks to variance assumed to be correlated with the shocks to the stock price.

To analyze the potential effects of stochastic volatility on the price of the autocall, recall from Section 3 that this structured product is best described as a combination of a long position in a plain vanilla bond and short positions in several digital options and in a put option, with the options having different maturities and the same underlying asset. The digital options, like the put option, pay if the underlying asset's price is low enough, and for all purposes behave similarly to put options. As described in Section 3, these options are out-of-the money at the time of issuance of the structured product, and may even be far-out-of-the money for a low enough threshold.

Consider a period of high volatility in which short-dated variance is above long-dated variance as suggested by the evidence in Section 2. Then the value of the options embedded in the autocall is lower relative to the model with constant volatility if volatility were always at the highest level. The magnitude of this effect depends on the persistence of shocks to volatility. When shocks to volatility last only one period, volatility quickly reverts down to its long-dated level, lowering the value of the options and raising the value of the autocall relative to the constant volatility model. When shocks to volatility are very persistent, then short-dated volatility can move away from its long-dated level for many periods. If the maturity of the autocall is sufficiently small, then the high persistence of variance of stock returns implies that the price of the autocall in the stochastic volatility model is close to the price in the constant volatility model of Section 4. Quantitatively, we show later in this section that these effects appear to be important determinants in the pricing of autocalls.

Heston (1993) shows that the volatility of volatility parameter controls the kurtosis of stock returns. Increasing volatility of volatility has the effect of increasing kurtosis, thus generating fatter tails in stock returns with the consequence that far-out-of-the money put option prices increase and near-the-money put option prices decrease, relative to the constant volatility model. Provided the threshold level on the autocall is sufficiently high, given the volatility in stock returns, the autocall represents a short position in near-the-money options and therefore we expect the price of the autocall to increase with a higher volatility of volatility, all else equal.

Heston (1993) shows that the correlation between shocks to the stock price and shocks to variance controls the skewness in stock returns. When this correlation is negative, stock returns display a "leverage effect" according to which low returns tend to be associated with high volatility and the distribution of stock returns is negatively skewed. Heston shows that then out-of-the money put option prices increase relative to the constant volatility model. We therefore expect the options in the autocall to increase in value when the correlation is negative. Therefore, the leverage effect is expected to lower the price of the autocall, all else equal.

We proceed to quantify the significance of these effects on the pricing of the Apple autocall that we have been studying. We let all parameters common to the model in Section 4 take on the same values and therefore assume that the short-dated volatility is 30%. In line with the evidence in Section 2, we consider two values for the unconditional volatility, 20% in the base case scenario and 30%. We calibrate the mean reversion parameter to match the mean value of persistence in asset volatility estimated in Engle and Siriwardane (2014). Following Engle and Siriwardane (2014), we set the mean reversion parameter to 0.4. Besides this base case, we also report results using a higher value of mean reversion of 3. When volatility reverts more rapidly to its mean it therefore displays less persistence. The magnitude of this parameter becomes critical when short-dated volatility and long-dated volatility differ significantly.

In the base case, we report results assuming no leverage effect, i.e. zero correlation. Because Apple's stock returns were negatively skewed at the date of issuance we also discuss cases with negative correlation. Finally, as we find that increasing volatility of volatility increases the value of the autocall, we calibrate the volatility of volatility to 12%. This value is the highest value that is consistent with the Feller condition that ensures that volatility is positive (see the discussion of the continuous-time version of the Heston model in Dragulescu and Yakovenko (2002)).

Table 3 presents the fair market value of the Apple autocall under several parameter combinations. The model is simulated using 50,000 price paths. In the baseline calibration, the fair market price is \$9.98, which implies a discount relative to the actual price of 0.2%. When we add a leverage effect by setting correlation to −0.2 we find that the fair market price drops to 9.95, representing an overpricing of 50 basis points. At the largest leverage effect, correlation of −1, the autocall overpricing is 1.3%. The relatively small overpricing in the base case calibration is due to the lower expected mean volatility, and to the presence of fat tails induced via volatility of volatility. Investors therefore appear to benefit from selling exposure to volatility risk to the underwriting institution.

[INSERT TABLE 3 HERE]

To understand the sources of gains to investors in this model relative to the model with constant volatility, we vary mean reversion and long-dated volatility. Increasing mean reversion in volatility produces a large increase in the value of the autocall for both values of the correlation. Intuitively, since long-dated volatility is lower than short-dated volatility, as would be expected in a scenario of relatively higher current volatility, the more rapidly volatility reverts back down to its long-dated level, the lower the value of the options embedded in the autocall and the higher the value of the autocall. Likewise, increasing long-dated volatility toward the short-dated level not only makes the value of mean reversion less relevant for the calibration, but also significantly lowers the value of the autocall because the value of the options in the autocall increases. But, the value of the autocall when long-dated volatility is 30% is only slightly higher than in the model with constant volatility indicating that the fat tails in the model produce a positive but small effect on the price of the autocall.

[INSERT TABLE 4 HERE]

To conclude, we discuss the patterns in conditional and unconditional probabilities that the security is called and the patterns in the simulated IRRs in the base case scenario. Table 4 presents the unconditional and conditional probabilities of the security being called in this model, also in the baseline case (from Table 2) and in the mean reversion case which will be discussed in the next section. There is a decline in the probability that the security is called after the first determination date due to the lower long-dated volatility, but this effect is not very large given the high persistence in volatility.

[INSERT FIGURE 6 HERE]

Figure 6 depicts the mean IRR conditional on survival for three models, the geometric Brownian motion of Section 4 (left plot) repeated from Figure 5A, the stochastic volatility model (center plot), and the model of mean reversion in price from Section 6 (right plot). For the stochastic volatility model, the unconditional mean IRR at the time of issue is 2.5%, higher than that in the constant volatility model or baseline case. This reflects the higher valuation and lower overpricing discussed above. The mean IRR conditional on survival declines as the expected capital loss at maturity increases. In the paths where the security is never called (right-most bars), the mean IRR is negative but less so than for the baseline model. This reflects the lower value of the options (including the put option at maturity) due to the mean reversion of volatility that appears to dominate the effect that volatility of volatility would have in generating a higher value for the options and hence a lower value for the autocall.

6. Price-level Effects on Autocall

The underlying asset price models studied in the previous sections have no role for price level effects. In this section we assume that the price of the underlying asset follows a mean reverting process, also known as an arithmetic Ornstein-Uhlenbeck process (see Dixit and Pindyck (1994)) (the Appendix contains the formal details). This price process allows for prices to fluctuate around a long-term mean, and to revert to that long-term mean at a fixed rate. To simulate this model, we calibrate the mean reversion parameter to 0.75, and the long-term mean of the price to 20% below the price at the issue date. According to this long-run mean, the seller of the autocall believes that there is a significant probability that the stock price will fall. In addition, noting that the volatility in this model is the volatility of the stock price and not the volatility of the stock return as in the previous geometric Brownian model, we adjust volatility to 0.317 so that the volatility of the stock price is matched in both models.

Simulating this model (also with 50,000 price paths) yields a fair market price of \$9.72, which implies a discount relative to the actual price of 2.8%, and an IRR computed using the expected cash flows across all simulated paths of 3.6% p.a., significantly lower than the required rate of 6.12%. The reason that the discount is so much larger relative to that under the baseline model of Section 4 (of 1.4%), and that the IRR is lower, is that the stock price at the issue date is significantly higher than its long run mean of \$8 which then has two main consequences: (i) the security is less likely to be called and to pay coupon plus interest at the first determination date; and, (ii) while it pays interest whenever the stock price is between \$7.5 and \$10, it is also more likely to drop below the threshold and result in a capital loss to the investor (note that the long run mean of the stock price is \$8 and the threshold is \$7.5). Table 4 contains the realized conditional and unconditional probabilities that the security is called. The unconditional probability of the security being called at the first determination date is only

37.2% and of being called in any of the first two determination dates is 47.6% (down from 64% in the baseline case). Also significant is that the unconditional probability that the security reaches maturity is 29% (up from 15% in the baseline case).

In Figure 6 we plot the mean IRR conditional on survival (right plot). The unconditional mean IRR at the time of issue (left-most bar) is 2.1% in the mean reversion model, higher than the 1.91% in the baseline model. Like in the baseline model, the mean IRR conditional on survival declines as the expected capital loss at maturity increases. This result explains the negative skewness in the IRR distribution and relies on the options embedded in the autocall.

Our results can explain the mixed evidence in Bergstresser (2008) on the effect of the recent past performance of the underlying asset on the likelihood of issuance. Whether issuance of a security is positively related with past performance of the underlying asset should depend on whether the structured security has embedded call-type options or put-type options. Because Bergstresser pools both call-like and put-like structured securities in his analysis of the likelihood of issuance, it is possible that the mixed evidence is caused by the lack of consideration of the separate and opposite effects of mean reversion in prices.

7. Conclusion

This paper describes the financial characteristics of a relatively new type of structured finance security, the autocallable contingent income security that has received significant attention because of the opportunity it gives investors to earn high coupons in a low yield environment. Yet, financial advisors that help sell these products often associate the high yield of these securities with their complexity and are careful in promoting these products to less sophisticated investors. These seemingly contradictory statements are the focus of this paper.

We offer two main takeaways. First, we document that underwriters of autocalls do not appear to choose underlying assets in a random fashion or to issue these securities at random times: the underlying security displays high volatility and is generally performing well in the stock market displaying prices at or near the 52-week high value. Second, we use this evidence and evidence from other financial markets to argue the inappropriateness of the workhorse model, the geometric Brownian motion, of the underlying asset's price used in the literature that values structured products. We show that when other models are considered that incorporate information from the underlying assets' characteristics at issuance of the note, specifically a model that allows for stochastic volatility, the valuation of autocalls appears fairly priced. These conclusions affect the valuation of structured products at large because we know from Bergstresser (2008) and Henderson and Pearson (2011) that our findings regarding the price characteristics of underlying assets at issuance apply also to other structured products. A broader study of price properties at the issuance of structured products is left for future research.

It is therefore possible that investors' views regarding price and volatility dynamics on the underlying assets may justify the overall interest in these securities and the growth in the market. Not-withstanding the finding that the appropriate choice of pricing model can approximate the fair value of the autocall to its actual price, private wealth managers and financial advisors in general should be aware of the potentially significant (remaining) overpricing due to the issuer's credit risk. We deliberately excluded consideration of this effect, but as the Lehman case suggests it can be large.

We leave one final remark from our analysis to financial analysts. Issuing banks engage in significant hedging of the exposures created by selling these structured products. Bennett and Gil (2012) alert for a potential "vicious circle" where a decrease in prices in equity markets associated with an increase in implied volatility can create a need to buy volatility as part of hedging by banks that in turn leads to an overshoot of volatility in a crisis. This effect may be particularly significant with underlying assets that have limited turnover in derivatives markets that can be used for hedging or at times when liquidity in these markets dries out (e.g. Millers (2013)).

References

1. Albuquerque, Rui. “Skewness in Stock Returns: Reconciling the Evidence on Firm versus Aggregate Stock Returns.” Review of Financial Studies 25.5 (May, 2012): 1630-1673.
2. Bennett, Colin, and Miguel A. Gil. “Volatility Trading.” (2012). Working paper Grupo Santander.
3. Bergstresser, Daniel. “The retail market for structured notes: Issuance patterns and performance, 1995-2008.” (2008). Working paper.
4. Bernard, Carole, Phelim P. Boyle, and William Gornall. “Locally Capped Investment Products and the Retail Investor.” *The Journal of Derivatives* 18.4 (Summer, 2011): 72-88.
5. Breuer, Wolfgang, and Achim Perst. “Retail banking and behavioral financial engineering: The case of structured products.” *Journal of Banking & Finance* 31.3 (2007): 827-844.
6. Burth, S, T. Kraus and H. Wohlwend. “The Pricing of Structured Products in the Swiss Market.” *Journal of Derivatives* 9 (2001): 30-40.
7. Célérier, Claire, and Boris Vallée. “What Drives Financial Complexity? A Look into the Retail Market for Structured Products.” (2014). Working paper.
8. Chang, Eric C., Dragon Yongjun Tang and Miao Zhang. “Suitability Checks and Household Investments in Structured Products.” *Journal of Financial and Quantitative Analysis* (forthcoming).
9. Financial Conduct Authority. “Structured Products.” http://www.fca.org.uk/firms/financial-services-products/investments/structured-products. (2013).
10. Coval, Joshua, Jakub Jurek and Erik Stafford. “The Economics of Structured Finance.” *The Journal of Economic Perspectives*, Vol. 23, No. 1 (Winter, 2009): 3-26.

11. Das, Sanjiv, Seoyoung Kim, and Meir Statman. "Liability Directed Investing in a Behavioral Portfolio Theory Framework." (2013). Working paper.
12. Deng, Geng, Tim Dulaney, Tim Husson, Craig McCann, and Mike Yan. "Ex-post Structured Product Returns: Index Methodology and Analysis." (2014). Working paper.
13. Deng, Geng, Guo Huali, and Craig McCann. " Structured Products in the Aftermath of Lehman Brothers." (2009). Working Paper.
14. Deng, Geng, Craig McCann, and Joshua Mallett. "Modeling Autocallable Structured Products." *Journal of Derivatives & Hedge Funds* 17.4 (2011): 326-340.
15. Dixit, Avinash K., and Robert S. Pindyck. "Investment under Uncertainty." Princeton University Press*, Princeton* (1994).
16. Dragulescu, A. A. , and V. M. Yakovenko. "Probability distribution of returns in the Heston model with stochastic volatility." *Quantitative Finance* 2 (2002): 443-453.
17. Engle, Robert F., and Emil N. Siriwardane. "Structural GARCH: The Volatility-Leverage Connection." (2014). Working paper.
18. Fabozzi, Frank J. "The Structured Finance Market: An Investor's Perspective." *Financial Analysts Journal* 61.3 (2005): 27-40.
19. Henderson, Brian J., and Neil D. Pearson. "The Dark Side of Financial Innovation: A Case Study of the Pricing of a Retail Financial Product." *Journal of Financial Economics* 100.2 (2011): 227-247.
20. Hens, Thorsten, and M. Rieger. "The Dark Side of the Moon: Structured Products from the Customer's Perspective." (2008). Working Paper.

21. Heston, Steven L. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *The Review of Financial Studies* 6.2 (1993): 327-343.

22. Kahl, C., and P. Jackel. "Fast Strong Approximation Monte-Carlo Schemes for Stochastic Volatility Models." *Quantitative Finance* 6.6 (2006): 513-536.

23. London Stock Exchange Group. "Structured Products." *http://www.londonstockexchange.com/specialist-issuers/structuredproducts/sp-introbrochure.pdf*. (2014). 5-6

24. Millers, Vita. "Korea Autocalls Pose Risk to Structured Products Market." *Risk.net* (2013): 4–7.

25. Oyo, Hiroshi. "SocGen: Japanese Uridashi Market Growing Despite Challenges." *http://www.structuredretailproducts.com/blog-story/17596/socgen-japanese-uridashi-market-growing-despite-challenges#.VHD-xtapq0M*, (2013).

26. Pereira da, Paulo, and Fernando Silva. "The Hidden Costs of Structured Retail Products Evidence for the Portuguese market." (2013). Working paper.

27. Purnell, Freya. "Why financial advisers are avoiding structured products." http://www.moneymanagement.com.au/analysis/financial-planning/2012/why-financial-advisers-avoid-structured-products, (2012).

28. Stein, J. *Overheating in Credit Markets: Origins, Measurement, and Policy Responses*, unpublished speech (2013) Federal Reserve Bank of St. Louis.

29. Stoimenov, Pavel A., and Sascha Wilkens. "Are Structured Products 'Fairly'Priced? An Analysis of the German Market for Equity-Linked Instruments." *Journal of Banking & Finance* 29.12 (2005): 2971-2993.

30. Szymanowska, Marta, Jenke Ter Horst, and Chris Veld. “Reverse Convertible Bonds Analyzed.” *Journal of Futures Markets* 29.10 (2009): 895-919.

31. Thin, Win. “The Carry Trade, Emerging Markets and Japanese Retail Investors.” *Credit Writedowns – Finance, Economics, Markets and Technology*, *https://www.creditwritedowns.com/2014/02/carry-trade-emerging-markets-japanese-retail-investors.html* (2014).

32. Wilkens, S., C. Erner, and K. Roder. “The Pricing of Structured Products in Germany.” *The Journal of Derivatives* 11 (Fall) (2003): 55–69.

33. Zweig, Jason. “How Apple’s Fall Bit Bondholders, Too.” *The Wall Street Journal* (Jan. 26, 2013).

Table 1: Sample Statistics

Panel A: Distribution by Year of Issuance

Issuance Year	*Principal Value*	*Number of Autocalls*
2009 (from 6/18)	$506,713,900	24
2010	$1,951,037,280	159
2011	$2,259,470,780	231
2012	$2,979,390,270	443
2013 (through 6/04)	$1,932,100,117	305
Total	*$9,628,712,347*	*1,162*

Panel B: Underwriters of Autocallable Securities

Underwriter	*Principal Value*	*Number of Autocalls*
JPMorgan Chase & Co.	$1,301,207,530	463
Morgan Stanley	$4,224,542,690	385
Citigroup Inc.	$2,874,140,487	172
Royal Bank of Canada	$477,715,980	57
UBS AG	$349,743,840	38
Barclays PLC	$284,412,680	24
Eksportfinans ASA	$26,100,000	10
HSBC USA Inc.	$81,580,960	9
Bank of America Corporation	$5,000,000	3
Credit Suisse AG	$4,264,580	1
Total	*$9,628,712,347*	*1,162*

Panel C: Categories of Autocallable Securities

Category	*Principal Value*	*Number of Autocalls*
1	$2,875,308,600	338
2	$264,926,000	162
3	$302,505,530	52
4	$144,399,610	30
5	$254,244,020	29
6	$14,685,000	28
7	$28,895,920	9
8	$88,421,000	5
9	$20,389,320	3
10	$2,785,000	1
11	$1,462,000	1
Multiple[a]		
2, 4, 11	$3,541,213,957	151
2, 7, 11	$237,050,000	116
4, 6, 11	$573,581,900	70
2, 4	$340,201,040	31
Total	*$8,742,655,997*	*1,026*

Notes: (a) Only multiple category combinations used in more than 30 issues are shown.

Categories: 1 – Standard autocall; 2 – Guaranteed coupon payment; 3 – Non-callable time span; 4 – No early redemption; 5 –Variable redemption level; 6 – No coupon payment; 7 –Multiple underlying assets; 8 –Maturity payment different; 9 – Variable coupon payment; 10 –Variable threshold level; 11 – Variable final payment

Panel D: Threshold Level Frequency Distribution

Threshold Level	*Frequency*	*Percent*	*Cumulative Percentage*	*Percent of Total Principal Value*
80%	296	25.47	25.47	38.48
75%	236	20.31	45.78	23.88
70%	215	18.50	64.29	15.92
65%	108	9.29	73.58	7.67
60%	102	8.78	82.36	3.85
90%	46	3.96	86.32	2.27
50%	39	3.36	89.67	2.42
85%	13	1.12	90.79	1.18
55%	12	1.03	91.82	0.20

Notes: Excludes autocalls with threshold levels observed fewer than 5 times.

Panel E: Coupon Rates in % per annum by Percentile and Coupon Rate Distribution Moments

E-1: Percentiles	*25%*	*50%*	*75%*	*Min*	*Max*	*Total Number*
Coupon rates	8.50	10.00	12.28	2.50	33.00	997

E-2: Distribution Moments	*Mean*	*Variance*	*Skewness*
	10.863	12.667	1.436

Note: Excludes autocalls with a variable coupon or no coupon payment.

Panel F: Autocall Maturity (in years) by Percentile and Maturity Distribution Moments

F-1: Years	*Principal Amount*	*Number of Autocalls*
< 1	$3,490,610,227	202
1	$2,792,102,730	504
1-4.99	$2,007,633,910	270
5	$496,067,480	73
5-9.99	$202,536,000	26
10	$85,794,000	16
10-14.99	$5,100,000	1
15	$393,471,000	46
18	$1,336,000	1
20	$154,061,000	25
Total	*$9,628,712,347*	*1,162*

F-2: Percentiles	*25%*	*50%*	*75%*	*Min*	*Max*	*Total Number*
Years	1	1	3	0.5	20	1,162

F-3: Distribution Moments	*Mean*	*Variance*	*Skewness*
	2.777	16.365	2.857

Panel G: Number of Autocalls Issued by Type of Underlying Asset

Type of Underlying Asset	*Principal Amount*	*Number of Autocalls*
Company	$7,233,029,787	748
Equity Index	$2,294,068,560	392
Commodity	$92,914,000	19
Currency	$8,700,000	3
Total	*$9,628,712,347*	*1,162*

Note: Of the autocalls that use companies as an underlying asset, 97.9% use a single company as the underlying asset with a principal of $7,184,776,267.

Panel H: Ratio of (Underlying Asset Implied Volatility/Underlying Asset Historical Implied Volatility) to (S&P Implied Volatility/S&P Historical Implied Volatility)

H-1: Percentiles	*25%*	*50%*	*75%*	*Min*	*Max*	*Total Number*
Ratio of volatilities	0.97	1.02	1.08	0.71	2.31	838

H-2: Distribution Moments	*Mean*	*Variance*	*Skewness*	*% obs. With Ratio>1*
	1.028	0.011	1.164	58.0

Note: There are 89 autocalls with the S&P 500 as the underlying asset, or 7.66%, of the data set. Bloomberg was used to obtain the 12-month implied volatility of both the underlying asset and the S&P implied volatility. The historical implied volatilities were obtained by using the respective historical average of the most recent 252-day implied volatilities. The total number of autocalls used in Panels H-1 and H-2 was reduced to 838 due to the elimination of those autocalls with multiple underlying assets, those having the S&P as the underlying asset, and those where the implied volatility was unavailable.

Panel I: Autocall Offering Price / 52-Week High with Single Underlying Asset

I-1: Percentiles	*25%*	*50%*	*75%*	*Min*	*Max*	*Total Number*
Offering price/52-week high	0.76	0.88	0.96	0.24	1	969

I-2: Distribution Moments	*Mean*	*Variance*	*Skewness*
	0.846	0.019	-1.183

Notes: Excludes autocalls with multiple underlying securities.

Panel J: Skewness of Log Returns

J-1: Percentiles	*25%*	*50%*	*75%*	*Min*	*Max*	*Total Number*
Skewness	-0.41	-.09	.16	-15.79	15.53	975

J-2: Distribution Moments	*Mean*	*Variance*	*Skewness*	*% obs. >0*
	-0.310	3.458	-3.230	57.5

Note: Excludes autocalls with multiple underlying assets and those where the pricing data was unavailable.

Table 2: Properties of Probabilities of Security Being Auto-Called

Model simulation of the probability that the security is called at any determination date. Baseline model is the Apple security presented in the text. Parameters in the baseline case are: risk free rate of 1.8% p.a., volatility of stock return of 30%, maturity of security of 3 years with quarterly determination dates, threshold level of 75% and coupon rate of 3.525% paid quarterly.

Notes: (a) Refers to the probability of reaching maturity. (b) Determination date 4 gives the probability of reaching maturity. (c) Probabilities after determination date 10 are omitted and determination 12 refers to maturity of the auto-call.

	Unconditional Probabilities											
Determination date	1	2	3	4	5	6	7	8	9	10	11	12(a)
Model												
Baseline case	0.5118	0.1286	0.0639	0.0397	0.0270	0.0209	0.0164	0.0128	0.0110	0.0098	0.0074	0.1507
Low volatility (15%)	0.5708	0.1352	0.0668	0.0411	0.0280	0.0213	0.0161	0.0126	0.0105	0.0084	0.0076	0.0816
High volatility (40%)	0.4916	0.1244	0.0606	0.0386	0.0281	0.0198	0.0168	0.0129	0.0113	0.0088	0.0077	0.1794
Low threshold (60%)	0.5118	0.1286	0.0639	0.0397	0.0270	0.0209	0.0164	0.0128	0.0110	0.0098	0.0074	0.1507
High threshold (85%)	0.5118	0.1286	0.0639	0.0397	0.0270	0.0209	0.0164	0.0128	0.0110	0.0098	0.0074	0.1507
Short maturity (1 year) (b)	0.5118	0.1286	0.0639	0.2957								
Long maturity (15 years) (c)	0.5118	0.1286	0.0639	0.0397	0.0270	0.0209	0.0164	0.0128	0.0110	0.0098	...	0.0549
Low coupon (8% p.a.)	0.5118	0.1286	0.0639	0.0397	0.0270	0.0209	0.0164	0.0128	0.0110	0.0098	0.0074	0.1507
High coupon (25% p.a.)	0.5118	0.1286	0.0639	0.0397	0.0270	0.0209	0.0164	0.0128	0.0110	0.0098	0.0074	0.1507

	Conditional Probabilities											
Determination date	1	2	3	4	5	6	7	8	9	10	11	12(a)
Model												
Baseline case	0.5118	0.2635	0.1778	0.1344	0.1055	0.0911	0.0789	0.0666	0.0615	0.0584	0.0466	0.0463
Low volatility (15%)	0.5708	0.3150	0.2273	0.1808	0.1507	0.1348	0.1177	0.1043	0.0970	0.0859	0.0850	0.0799
High volatility (40%)	0.4916	0.2447	0.1577	0.1193	0.0988	0.0770	0.0711	0.0585	0.0547	0.0447	0.0414	0.0365
Low threshold (60%)	0.5118	0.2635	0.1778	0.1344	0.1055	0.0911	0.0789	0.0666	0.0615	0.0584	0.0466	0.0463
High threshold (85%)	0.5118	0.2635	0.1778	0.1344	0.1055	0.0911	0.0789	0.0666	0.0615	0.0584	0.0466	0.0463
Short maturity (1 year) (b)	0.5118	0.2635	0.1778	0.1296								
Long maturity (15 years) (c)	0.5118	0.2635	0.1778	0.1344	0.1055	0.0911	0.0789	0.0666	0.0615	0.0584	...	0.0120
Low coupon (8% p.a.)	0.5118	0.2635	0.1778	0.1344	0.1055	0.0911	0.0789	0.0666	0.0615	0.0584	0.0466	0.0463
High coupon (25% p.a.)	0.5118	0.2635	0.1778	0.1344	0.1055	0.0911	0.0789	0.0666	0.0615	0.0584	0.0466	0.0463

Table 3: Autocall Prices in the Stochastic Volatility Model

Unless otherwise noted model parameters are: initial volatility of stock price of 30%, long-run mean volatility ($\bar{v}$) of 20%, mean reversion (λ) of 0.4, correlation between stock price shocks and volatility shocks (ρ) of 0, and volatility of variance (θ) of 12%.

	θ=0.12	
	λ=0.4	λ=3
$\bar{v}^{1/2} = 0.2$		
ρ=0	9.98	10.17
ρ=-0.5	9.93	10.15
$\bar{v}^{1/2} = 0.3$		
ρ=0	9.87	9.89
ρ=-0.5	9.85	9.86

Table 4: Properties of Probabilities of Security Being Auto-Called

Model simulation of the probability that the security is called at any determination date. Baseline model parameters are as indicated in Table 2.
Stochastic volatility model parameters: initial volatility of stock price of 30%, long-run mean volatility of 20%, mean reversion of 0.4, correlation between stock price shocks and volatility shocks of 0, and volatility of variance of 12%.
Mean reversion model parameters: volatility of stock price of 31.7%, long run mean price 20% below price of underlying at issuance, and parameter controlling the speed of adjustment equals 0.75.

Notes: (a) Refers to the probability of reaching maturity.

	Unconditional Probabilities											
Determination date	1	2	3	4	5	6	7	8	9	10	11	12(a)
Model												
Stochastic Volatility	0.5129	0.1242	0.0619	0.0386	0.0271	0.0194	0.0167	0.0121	0.0109	0.0095	0.0080	0.1588
Mean Reversion	0.3719	0.1043	0.0568	0.0379	0.0298	0.0245	0.0216	0.0181	0.0170	0.0150	0.0131	0.2900
Baseline case	0.5118	0.1286	0.0639	0.0397	0.0270	0.0209	0.0164	0.0128	0.0110	0.0098	0.0074	0.1507
	Conditional Probabilities											
Determination date	1	2	3	4	5	6	7	8	9	10	11	12(a)
Model												
Stochastic Volatility	0.5129	0.2550	0.1706	0.1284	0.1031	0.0823	0.0772	0.0607	0.0580	0.0537	0.0482	0.0467
Mean Reversion case	0.3719	0.1661	0.1084	0.0812	0.0694	0.0615	0.0576	0.0513	0.0508	0.0472	0.0431	0.0435
Baseline case	0.5118	0.2635	0.1778	0.1344	0.1055	0.0911	0.0789	0.0666	0.0615	0.0584	0.0466	0.0463

Figure 1: Coupon Rate Distribution

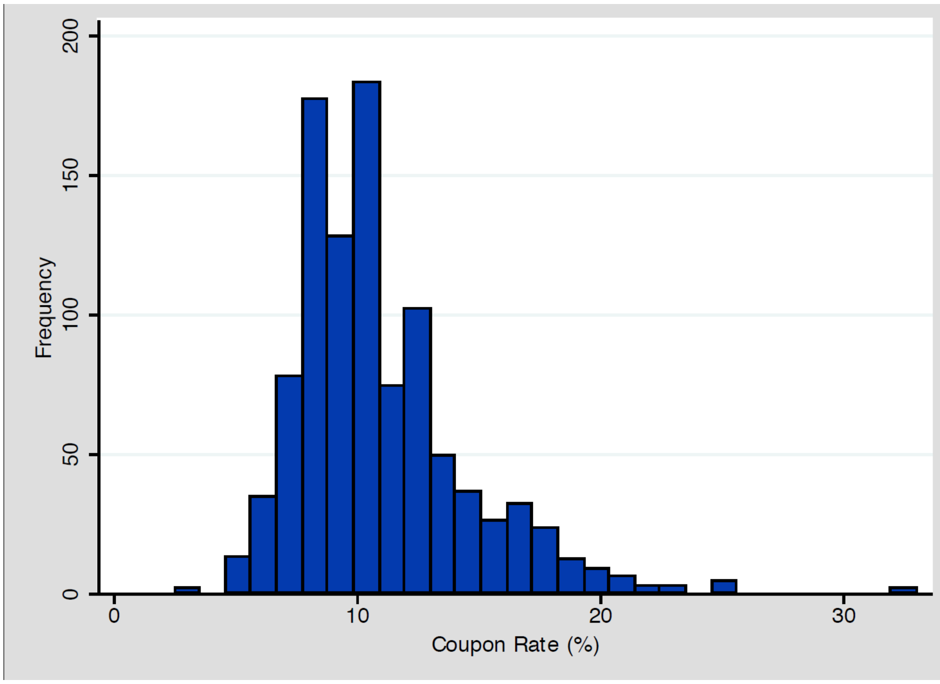

Note: Excludes autocalls without a constant coupon payment.

Figure 2: Top 10 Most Frequent Underlying Assets for Autocall Securities

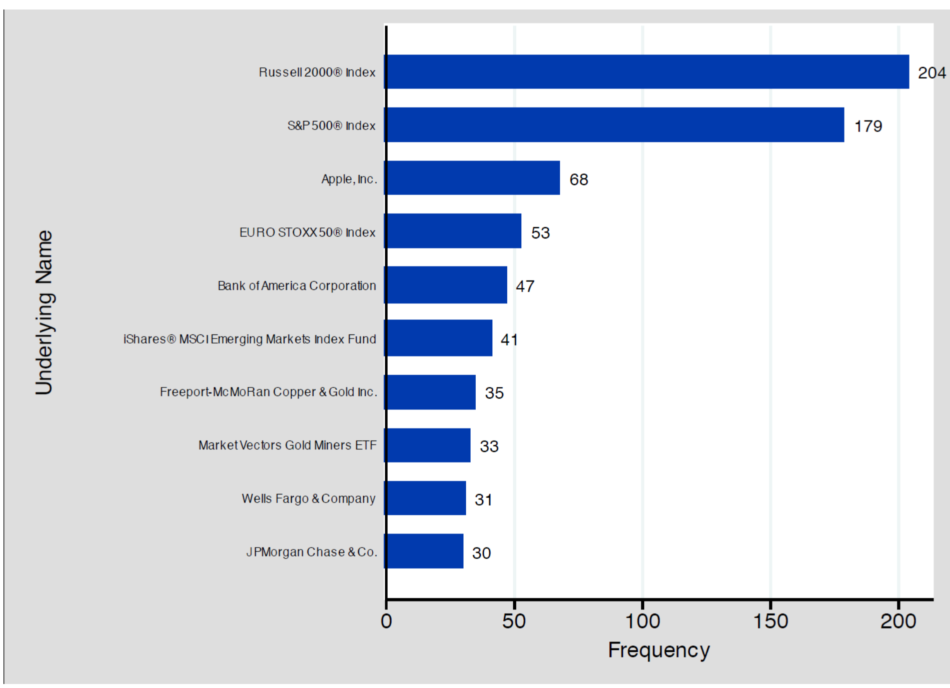

Note: iShares® Russell 2000 Index Fund is categorized together with Russell 2000 Index.

Figure 3: Ratio of (Underlying Asset's Implied Volatility / Underlying Asset's Historical Implied Volatility) to (S&P's Implied Volatility / S&P's Historical Implied Volatility)

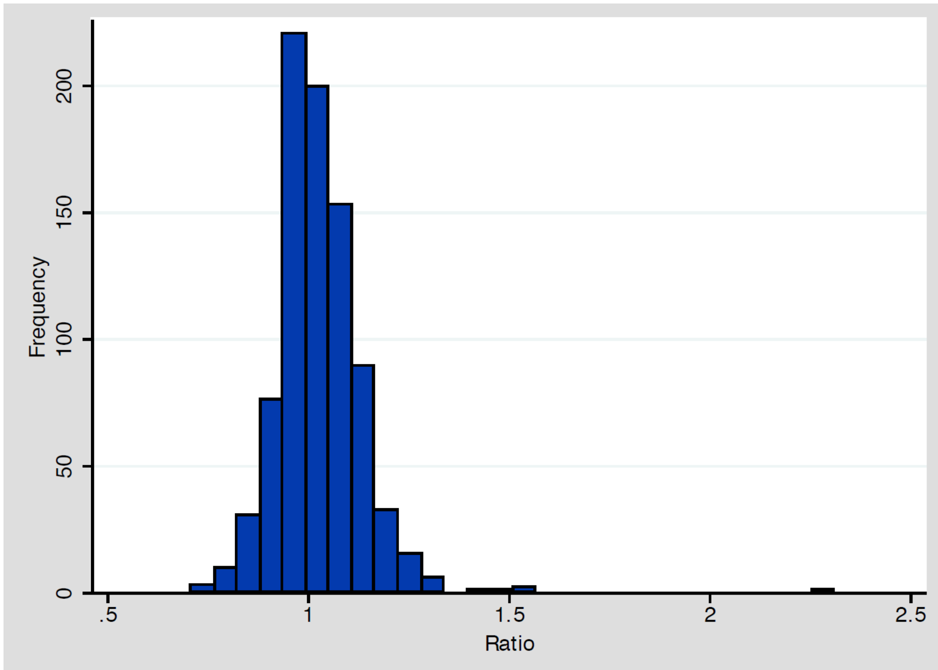

Note: Excludes autocalls with multiple underlying assets, those having the S&P as an underlying asset, and those where the implied volatility was unavailable.

Figure 4: Ratio of Initial Price of Underlying Asset to 52-Week High

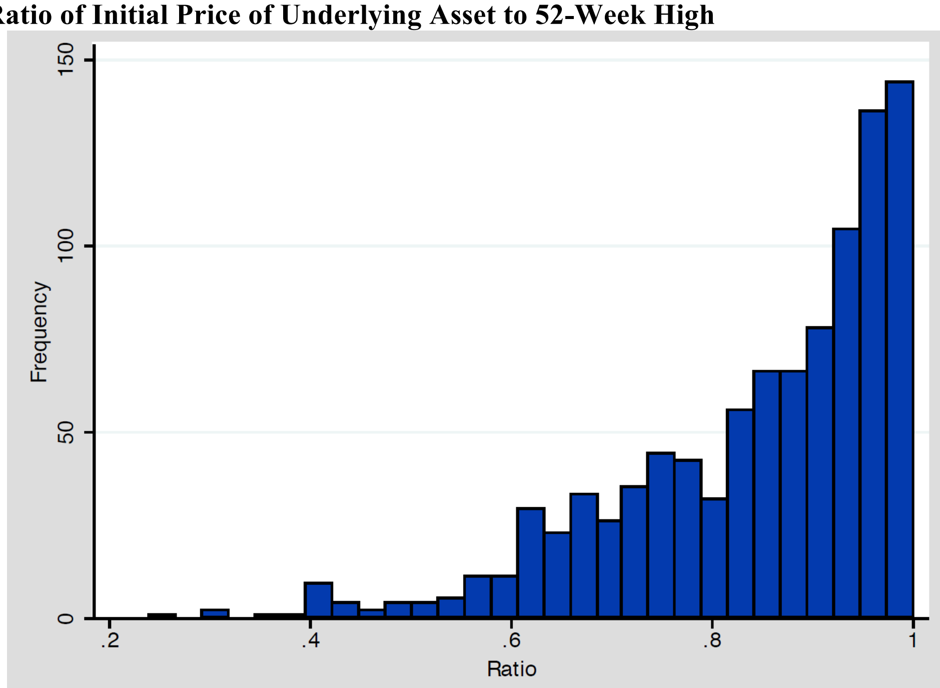

Note: Excludes autocalls with multiple underlying assets.

Figure 5A: Simulated IRRs of the Contingent Income Autocallable Security: Baseline Model and Low and High Volatility Models

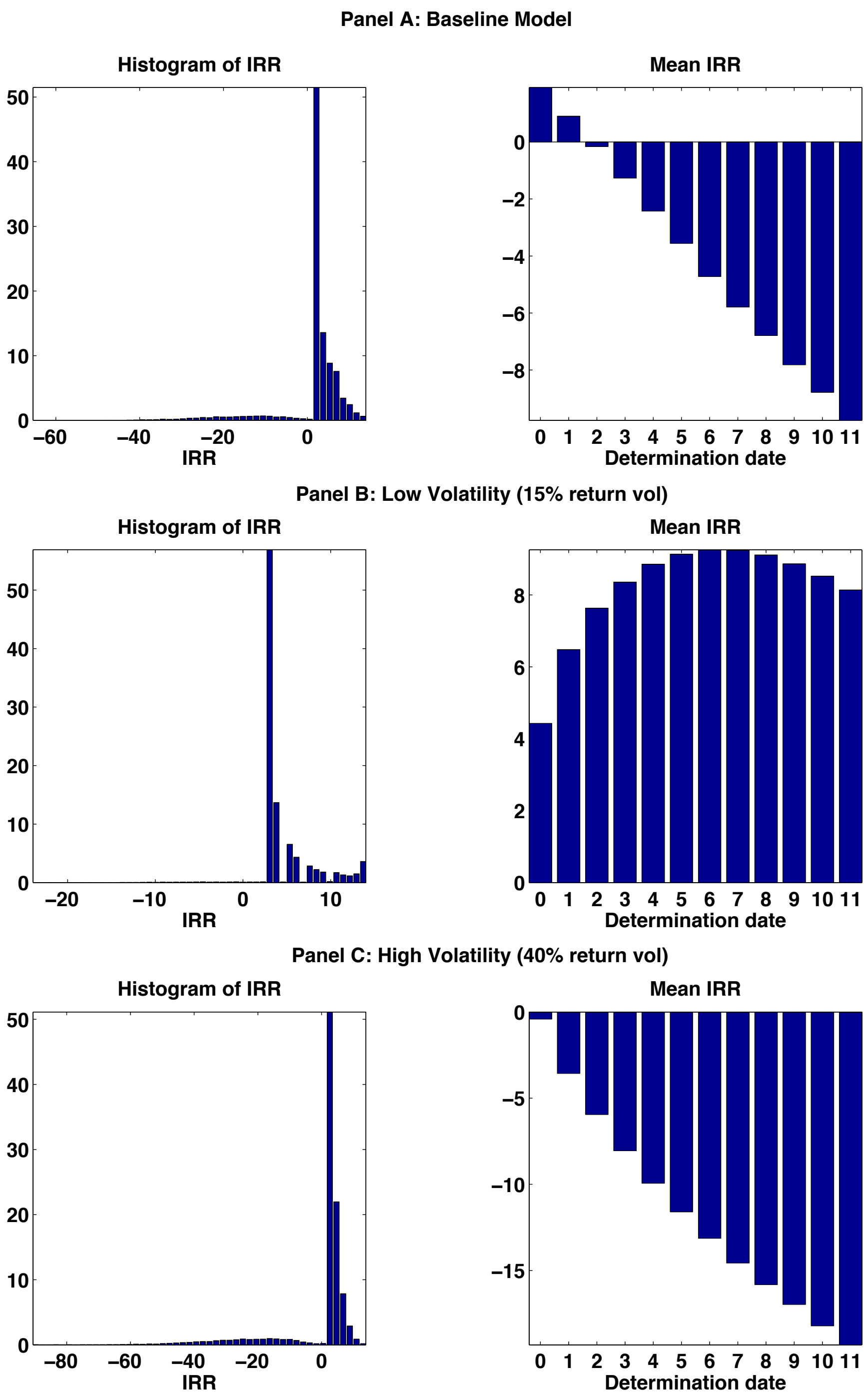

Note: The six plots are organized in the following manner. The first column contains plots labeled "Histogram of IRR" that depict the frequency distribution (in %) of the annualized IRR across all simulated paths. The second column contains plots labeled "Mean IRR" that depict the average annualized IRR (in %) computed across all simulated paths that survive a given determination date.

Figure 5B: Simulated IRRs of the Contingent Income Autocallable Security: Baseline Model and Low and High Threshold Models

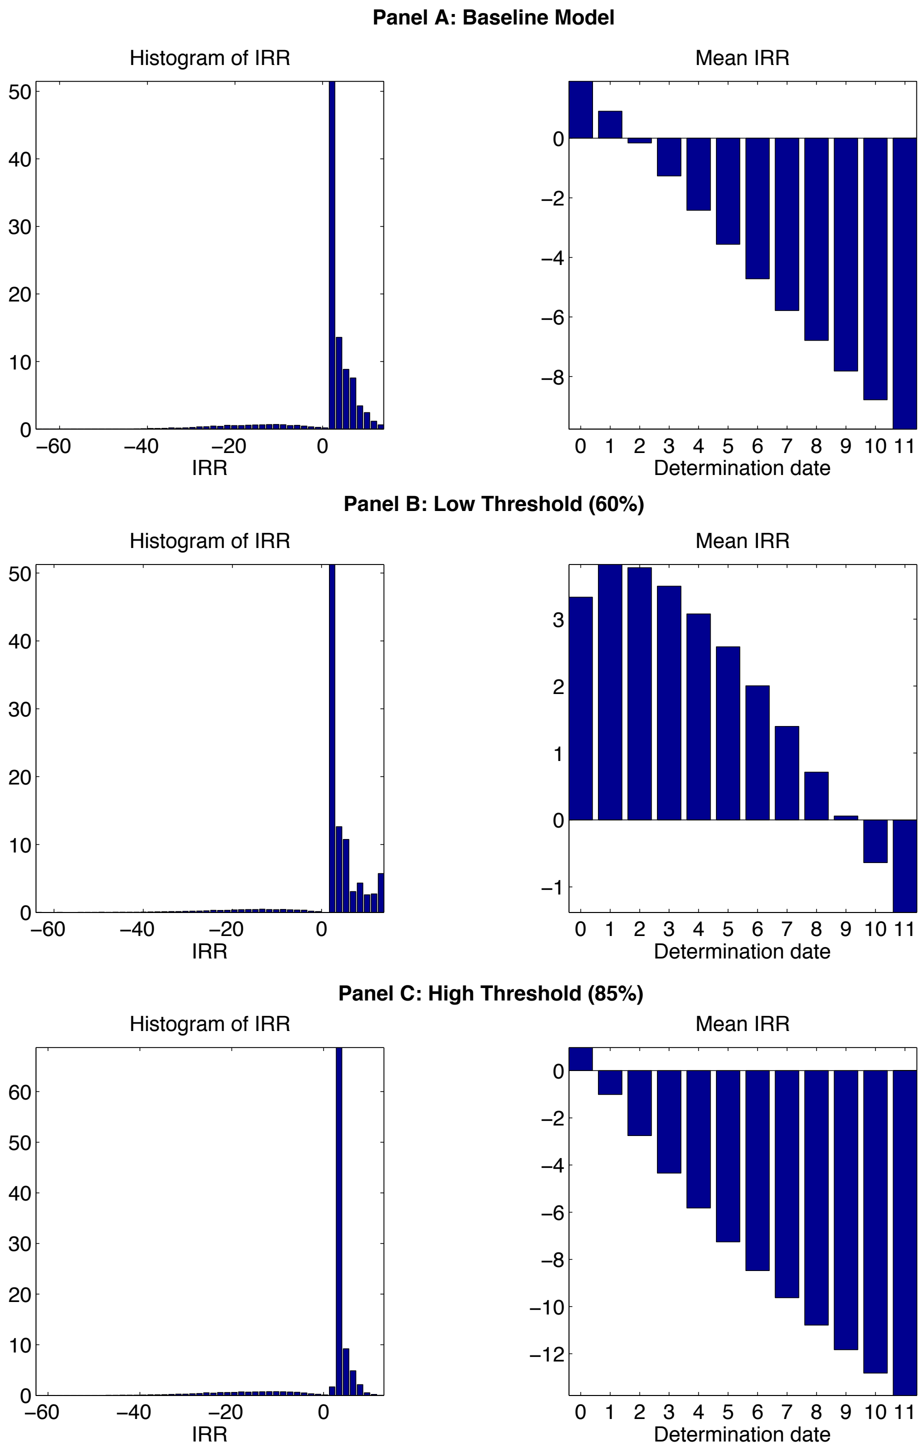

Note: Refer to notes in Figure 5A.

Figure 5C: Simulated IRRs of the Contingent Income Autocallable Security: Baseline Model and Short and Long Maturity Models

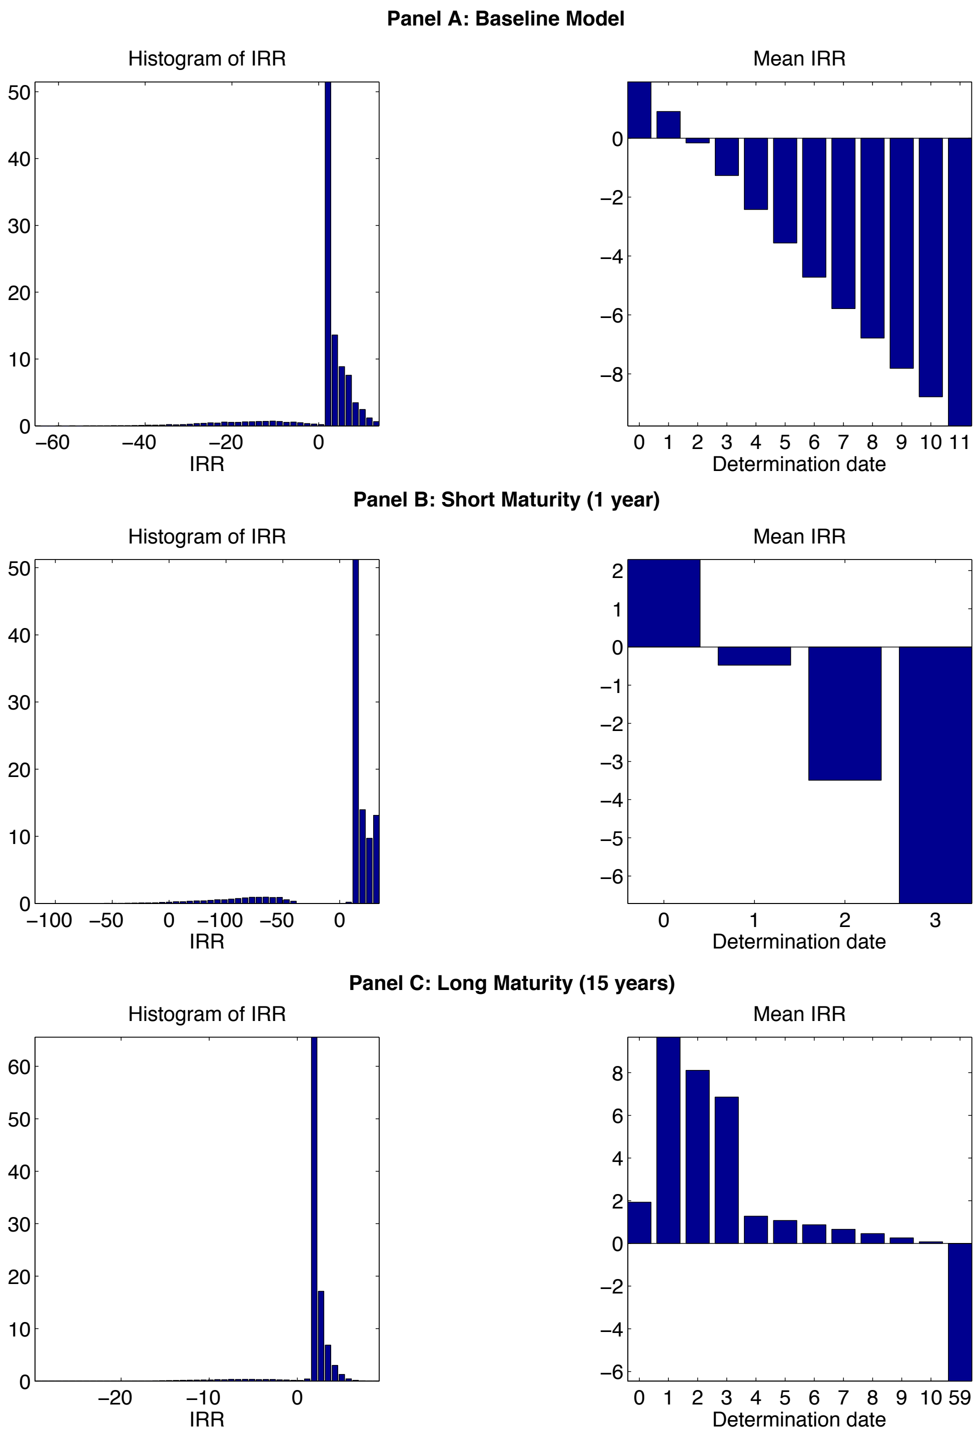

Note: Refer to notes in Figure 5A. In the corner, right-most plot determination dates jump from date 10 to the last determination date, date 59, to make the data easier to visualize.

Figure 5D: Simulated IRRs of the Contingent Income Autocallable Security: Baseline Model and Low and High Coupon Models

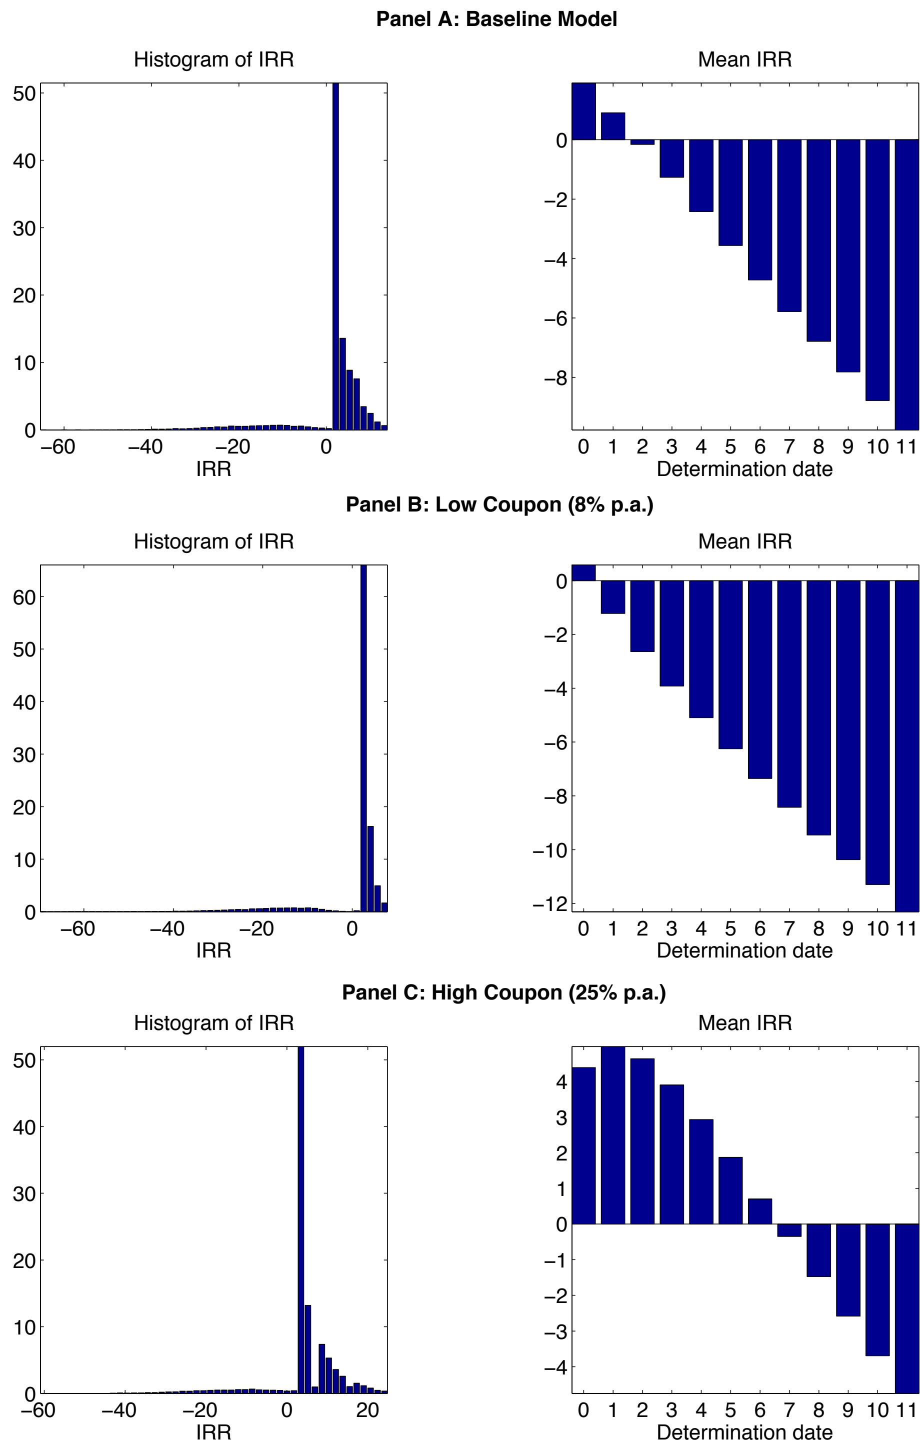

Note: Refer to notes in Figure 5A.

Figure 6: Simulated IRRs Under Different Underlying Asset Pricing Models

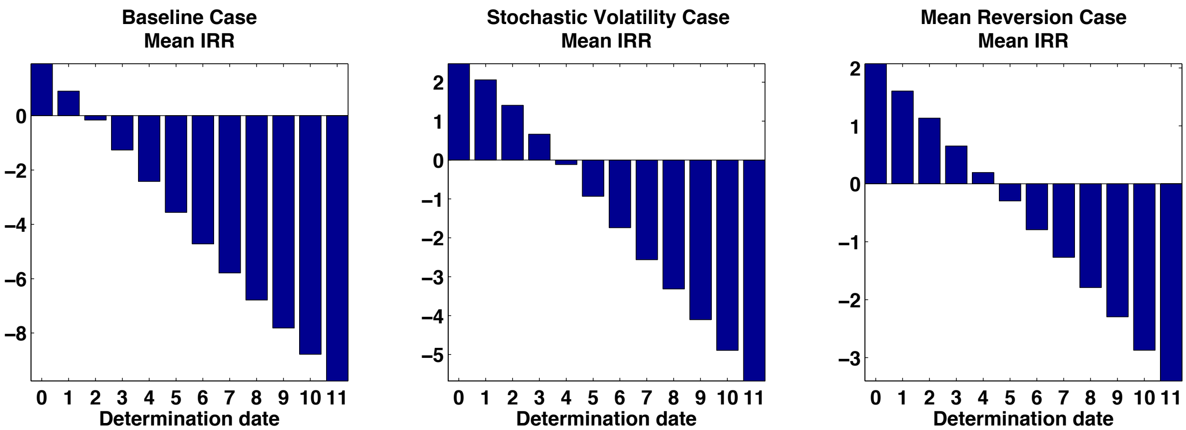

Note: Mean annualized IRRs by determination date in %.

Appendix

Framework for Analysis

The contingent income autocallable security that we analyze is of the plain vanilla kind and has the following features:

- Let t denote time, $t = 0$ be the issue date and $t = T$ be the autocall's maturity;
- Payouts are a function of the price performance of an underlying asset; the price of the underlying at the issue date is S_0; for simplicity the price and notional value of the autocall is $P = S_0$;
- At determination date $t = 1, \dots, T-1$, the security is called if $S_t > S_0$, in which case the investor gets $(1+i) \times P$, where i is the coupon rate, and no further cash flows, or is not called, in which case the investor gets $i \times P$ if $S_t > \alpha S_0$, with $\alpha < 1$ (αS_0 is the threshold level);
- At maturity, the autocall pays $(1+i) \times P$ if $S_T > \alpha S_0$, or otherwise it pays either one unit of the underlying asset or its current cash value of S_T. Clearly, if the underlying asset is received, the investor has a capital loss of $S_0 - S_T > (1-\alpha)S_0$.

We study pricing of the autocall under three models that describe the price behavior of the underlying asset. The first model assumes that the price of the underlying asset, S_t, follows a geometric Brownian motion:

$$\frac{S_t}{S_0} = exp\left\{\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma W_t\right\}, \quad (1)$$

where μ is the instantaneous growth rate in prices, σ is the instantaneous return volatility, and W_t is a Wiener process whose continuous increments are normally distributed with mean zero and unit variance. To simulate the process we discretize the process in Equation (1) using[13]

[13] Deng et al. (2014) provide an approximate analytical solution to the price of an autocall.

$$\frac{S_{t+\Delta t}}{S_t} = exp\left\{\left(\mu - \frac{1}{2}\sigma^2\right)\Delta t + \sigma N(0,1)\sqrt{\Delta t}\right\}. \quad (2)$$

An important assumption of the geometric Brownian motion model is that of constant volatility. The second model we study relaxes this assumption. We let the price process follow the Heston (1993) model, which in discretized form is

$$\frac{S_{t+\Delta t}}{S_t} = exp\left\{\left(\mu - \frac{1}{2}\mathrm{v}_t\right)\Delta t + W_t\sqrt{\mathrm{v}_t\Delta t}\right\},$$

and the instantaneous variance of the stock return, v_t, follows the process[14]

$$\mathrm{v}_{t+\Delta t} = \mathrm{v}_t + \lambda(\bar{\mathrm{v}} - \mathrm{v}_t)\Delta t + \theta Z_t\sqrt{\mathrm{v}_t\Delta t} + \frac{1}{4}\theta^2\Delta t\left(Z_t{}^2 - 1\right). \quad (3)$$

In this model λ dictates the speed of mean reversion in variance and θ is the instantaneous volatility of variance. The standard normal shocks in the price equation and the variance equation, respectively, W_t and Z_t, are assumed correlated with a correlation coefficient of ρ. v_t is the conditional variance of stock returns, which we label as short-dated variance. $\bar{\mathrm{v}}$ is the long-term mean of the conditional variance, which we label as long-dated variance.

The third and last model we study allows for price level effects. We let the price of the underlying asset follow a mean reverting process, also known as an arithmetic Ornstein-Uhlenbeck process (see Dixit and Pindyck (1994)):

$$S_{t+\Delta t} = \exp\left\{\begin{matrix} \ln(S_t)\exp(-\eta\Delta t) + \ln(\bar{S})\,(1 - \exp(-\eta\Delta t)) \\ -(1 - \exp(-2\eta\Delta t))\frac{\sigma^2}{4\eta} + \sigma\sqrt{\frac{1-\exp(-2\eta\Delta t)}{2\eta}}N(0,1) \end{matrix}\right\}. \quad (4)$$

In this formula, η is the parameter that controls the speed of mean reversion (i.e., $\ln(2)/\eta$ is the half life of a shock to prices), and $\bar{S}$ is the long run mean of the price.

[14] The discretized variance equation (Equation (3)) in the Heston model does not guarantee that variance is nonnegative. Several approaches have been proposed to minimize this concern. In this paper, we use the Milstein scheme (see Kahl and Jackel (2006)). The Milstein scheme adds the last term on the right-hand side of the variance equation. We simulate the model with various discretizations with similar results across approaches.