Financial Analysts Journal
Volume 71 • Number 3
©2015 CFA Institute

Investment Analysis of Autocallable Contingent Income Securities

Rui Albuquerque, Raquel M. Gaspar, and Allen Michel

Autocallable contingent income securities (autocalls) have payouts contingent on the performance of an underlying asset and give investors an opportunity to earn high yields in a low-interest environment. The authors collected data on US-issued autocalls and modeled a typical autocall under various assumptions, finding that they are issued on underlying assets that display high volatility, high prices, and negative skewness. Incorporating stochastic volatility into the model explains some of the overpricing routinely reported in prior studies.

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

This growth notwithstanding, financial advisers who help intermediate the sale of these products often associate their high yields with their complexity and are reluctant to promote them outside of their high-net-worth client pool and those sophisticated investors who have a view on the market (see Purnell 2012). Recent investment regulation in Europe (the European Commission's Markets in Financial Instruments Directive, or MiFID II), the United States (the Dodd–Frank Act), the United Kingdom (see Financial Conduct Authority 2013), and other countries (see references in Chang, Tang, and Zhang, forthcoming) has pushed for greater investor protection, including the separation of independent and nonindependent advice, limitations on the receipt of commissions, and even the prohibition of marketing and distribution of certain products (suitability requirements).

The purpose of this article is to further our understanding of these seemingly contradictory features of the marketplace for structured products by presenting an analysis of their pricing and return characteristics. There are two main challenges in studying the pricing of structured notes. First, these products are extremely complex and have 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 (see Bergstresser 2008; Célérier and Vallée 2014) cannot distill the causes for the observed pricing from the properties of the notes. In this article, 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 US market and are part of a body of structured finance offerings called "reverse convertibles," whose payoffs are 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 16 July 2012 autocall with Apple Inc. as the underlying asset.[1] Apple stock on 16 July 2012 opened at $605.12, within 7% of its prior all-time high of $644, reached on 10 April 2012. With a maturity of three years, the autocall offered a coupon of 3.525% per quarter (approximately 14% per year),[2] significantly higher than the rate available in the marketplace at that time. The investor receives the coupon only if

Rui Albuquerque is associate professor of finance at Boston University School of Management and a full professor at Católica-Lisbon School of Business and Economics, Portugal. Raquel M. Gaspar is a researcher at ISEG–CEMAPRE, Universidade de Lisboa, Lisbon, Portugal. Allen Michel is professor of finance at Boston University School of Management.

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

We hand collected data on all autocallable contingent income securities listed in the SEC's EDGAR database from June 2009 to June 2013, totaling $9.6 billion of notional value. In terms of product contractual characteristics, we found that two-thirds of all autocalls differed in at least one of the many product features as compared with the plain vanilla example of the Apple autocall.[3] We also found that the median annualized coupon rate is 10% and the median maturity is one year.

In terms of characteristics of the underlying asset, we found that for 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 one-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 S&P 500 Index volatility; 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 (these last two results are consistent with the evidence in Henderson and Pearson [2011] for SPARQS [Stock Participation Accreting Redemption Quarterly-Pay Securities], 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.

We valued the autocall under three alternative price processes for the underlying asset. We started by assuming that the underlying asset's price follows a geometric Brownian motion, which is the most common model in the literature. We found that the Apple autocall described earlier has a fair market price of $9.86, or 1.4% below the actual price. We then departed from the geometric Brownian motion model in evaluating autocalls—and thus departed from the vast majority of the literature—and studied a model of stochastic volatility and a model of mean reversion in prices. We believe that geometric Brownian motion is not an appropriate model for the underlying assets' price because of several factors. These factors include (1) our empirical findings, those of Henderson and Pearson (2011), those of Bergstresser (2008), and practitioner discussions (see Millers 2013) regarding the underlying assets' price characteristics at issuance that suggest that underwriters do not choose underlying assets at random; (2) the large body of evidence of stochastic volatility showing differences in short- and long-term volatility; and (3) 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-term volatility) on the valuation of autocalls can differ from long-term volatility, which is not possible under the geometric Brownian motion model, in which 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 underwriters of underlying assets that trade at high prices at issuance. This is not possible under the geometric Brownian motion model, in which future returns do not depend on the initial price. We found that when prices are allowed to mean revert, 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 costs for 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 and is 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. In a study of the Portuguese structured retail product market, Pereira da Silva and Silva (2013) found that this hidden credit cost averages 4.9% per year. Deng, Huali, and McCann (2009) found that with the increased borrowing costs faced by Lehman Brothers 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 et al. 2011). The other feature is the potential lack of liquidity of these assets. In the United States, 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 market making by the issuer or the London Stock Exchange (see London Stock Exchange 2014). Overall, the evidence in our article 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 have examined these properties. Bergstresser (2008) studied a vast array of notes with call- and put-like options embedded, and Henderson and Pearson (2011), like us, studied 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 (see Burth, Kraus, and Wohlwend 2001; Henderson and Pearson 2011; Stoimenov and Wilkens 2005). To our knowledge, only two papers besides ours have deviated from this benchmark. Pereira da Silva and Silva (2013) also used a Heston model, and Célérier and Vallée (2014) used a local volatility model, but neither paper's model was calibrated to reflect the discrepancy between long-term volatility and volatility at the issuance date that we and others have found.[4] The other approach values the structured product by replicating its payoff using bonds and options traded (see Burth et al. 2001; Wilkens, Erner, and Roder 2003). The advantage of this approach is that it is "model free." The disadvantage is the difficulty in accounting for transaction costs necessary for replication.

Using geometric Brownian motion as a model for the underlying asset's price, the literature generally has found that underwriters overprice the structured products they sell (see Wilkens et al. 2003; Bergstresser 2008; Henderson and Pearson 2011; Bernard, Boyle, and Gornall 2011; Deng, Dulaney, Husson, McCann, and Yan 2014). Whereas we point to an explanation for overpricing that has to 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 transaction cost explanations that rely on the inability of retail investors to trade in certain markets (e.g., futures and option markets) at the same prices that large institutions can. The general sense is that these are not large enough to explain the findings (see Bergstresser 2008; Henderson and Pearson 2011). Another common clientele explanation for the demand for high-yield structured products is the low-rate environment and the ability to achieve some degree of capital protection (see Burth et al. 2001; Stoimenov and Wilkens 2005; Coval, Jurek, and Stafford 2009; Szymanowska et al. 2009; Stein 2013).

Several behavioral explanations have been advanced, including investor irrationality or bounded rationality, framing, and overweighting of low-probability events (see Breuer and Perst 2007; Bergstresser 2008; Hens and Rieger 2008; Bernard et al. 2011; Henderson and Pearson 2011; Das, Kim, and Statman 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 (see Stoimenov and Wilkens 2005; Pereira da Silva and Silva 2013; Célérier and Vallée 2014), but it is less clear why it would predict overpricing.

Autocall Sample Characteristics

In this section, we discuss the sample characteristics, including both those that are unrelated and those that are related to the underlying asset. Such an assessment is important for understanding the nature of the securities on which autocalls are written.

Characteristics Unrelated to the Underlying Asset. We collected the universe of contingent income autocallable securities from the SEC's EDGAR database from 18 June 2009 through 4 June 2013. We searched 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 \$9 billion.[5] Panel B of Table 1 lists the underwriters. There are two main bank underwriters, JPMorgan

Table 1. Sample Statistics

	Principal Value	Number of Autocalls
A. Distribution by year of issuance		
2009 (starting 18 June)	$506,713,900	24
2010	1,951,037,280	159
2011	2,259,470,780	231
2012	2,979,390,270	443
2013 (through 4 June)	1,932,100,117	305
Total	$9,628,712,347	1,162
B. Underwriters of autocallable securities		
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
C. Categories of autocallable securities		
1. Standard autocall	$2,875,308,600	338
2. Guaranteed coupon payment	264,926,000	162
3. Noncallable time span	302,505,530	52
4. No early redemption	144,399,610	30
5. Variable redemption level	254,244,020	29
6. No coupon payment	14,685,000	28
7. Multiple underlying assets	28,895,920	9
8. Maturity payment different	88,421,000	5
9. Variable coupon payment	20,389,320	3
10. Variable threshold level	2,785,000	1
11. Variable final payment	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

D. Threshold-level frequency distribution (%)[b]

Threshold Level	Frequency	Percentage	Cumulative Percentage	Percentage 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

(continued)

Table 1. Sample Statistics (continued)

E. Coupon rates in % per year by percentile and coupon rate distribution moments[c]

	Percentile					
	25%	50%	75%	Min.	Max.	Total No.
Coupon rate	8.50	10.00	12.28	2.50	33.00	997

Distribution moment	Mean	Variance	Skewness
	10.863	12.667	1.436

F. Autocall maturity (in years) by percentile and maturity distribution moments

Years	Principal Value	Number of Autocalls
< 1	$3,490,610,227	202
1	2,792,102,730	504
1.01–4.99	2,007,633,910	270
5	496,067,480	73
5.01–9.99	202,536,000	26
10	85,794,000	16
10.01–14.99	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

	Percentile					
	25%	50%	75%	Min.	Max.	Total No.
Years	1	1	3	0.5	20	1,162

Distribution moment	Mean	Variance	Skewness
	2.777	16.365	2.857

G. Number of autocalls issued by type of underlying asset[d]

Type of Underlying Asset	Principal Value	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

H. Ratio of (underlying asset implied volatility/underlying asset historical implied volatility) to (S&P 500 implied volatility/S&P 500 historical implied volatility)[e]

	Percentile					
	25%	50%	75%	Min.	Max.	Total No.
Ratio of volatilities	0.97	1.02	1.08	0.71	2.31	838

Distribution moment	Mean	Variance	Skewness	Obs. with Ratio > 1
	1.028	0.011	1.164	58.0%

(continued)

Table 1. Sample Statistics (continued)

I. Autocall offering price/52-week high with single underlying asset[f]

	Percentile					
	25%	50%	75%	Min.	Max.	Total No.
Offering price/52-week high	0.76	0.88	0.96	0.24	1	969
Distribution moment	Mean	Variance	Skewness			
	0.846	0.019	–1.183			

J. Skewness of log returns[g]

	Percentile					
	25%	50%	75%	Min.	Max.	Total No.
Skewness	–0.41	–0.09	0.16	–15.79	15.53	975
Distribution moment	Mean	Variance	Skewness	Obs. > 0		
	–0.310	3.458	–3.230	57.5%		

[a]Only multiple category combinations used in more than 30 issues are shown.
[b]Excludes autocalls with threshold levels observed fewer than five times.
[c]Excludes autocalls with a variable coupon or no coupon payment.
[d]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.
[e]There are 89 autocalls with the S&P 500 as the underlying asset, representing 7.66% of the dataset. Bloomberg was used to obtain the 12-month implied volatility of both the underlying asset and the S&P 500. 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 Panel H was reduced to 838 owing to the elimination of those autocalls with multiple underlying assets, those having the S&P 500 as the underlying asset, and those for which the implied volatility was unavailable.
[f]Excludes autocalls with multiple underlying securities.
[g]Excludes autocalls with multiple underlying assets and those for which the pricing data were unavailable.

Chase and Morgan Stanley, representing 73% of all issuances and 57% of the issued dollar volume.

There are many variations on the plain vanilla autocall described in the previous section. 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 or below the threshold level, and has the possibility of early redemption. Panel C of Table 1 shows that roughly two-thirds of all autocalls differ from the vanilla autocall in at least one feature.

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 year basis, by percentile, as well as their distribution moments. Because 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% per year 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% per year. Panel F of Table 1 shows that both the mode and the median maturity are 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, representing 23.2% of the autocalls in the sample.

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), representing 64.7% of all issues, or 75% of the principal amount underwritten. The next most frequently underwritten securities are equity indexes and commodities. **Figure 2** shows that the most common underlying assets are the Russell 2000 Index (17.6%), the S&P 500 Index (15.4%), and Apple (5.9%).

Panel H presents statistics computed on the basis of the stock's volatility relative to its historical volatility. We calculated the following ratio: (underlying asset's implied volatility/underlying asset's historical implied volatility)/(S&P 500 implied volatility/S&P 500 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 found that in about 60% of

Figure 1. Coupon Rate Distribution

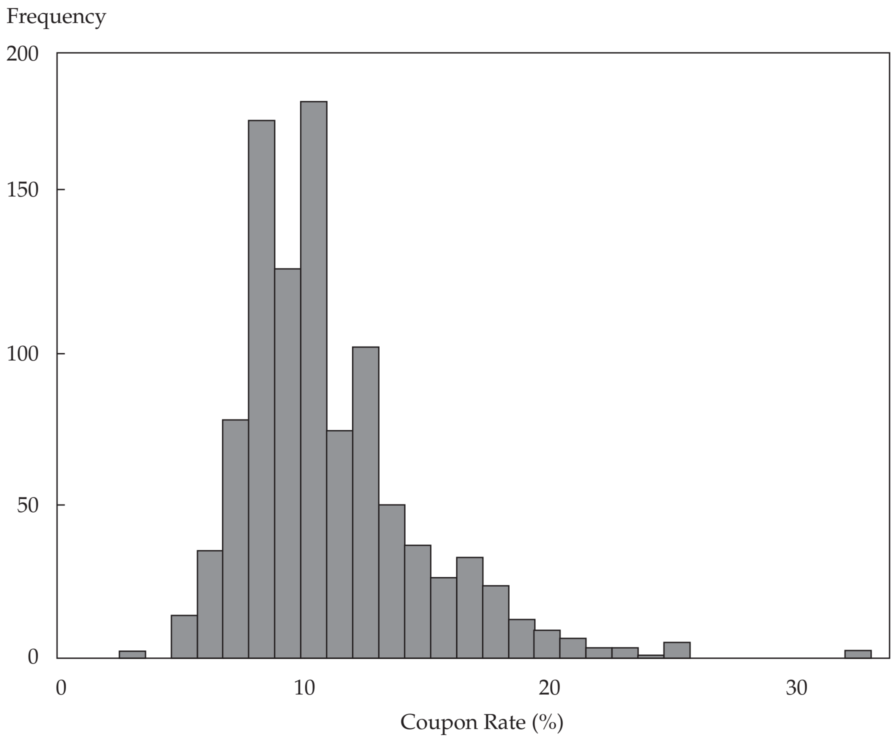

Note: Excludes autocalls without a constant coupon payment.

Figure 2. Top 10 Most Common Underlying Assets for Autocall Securities

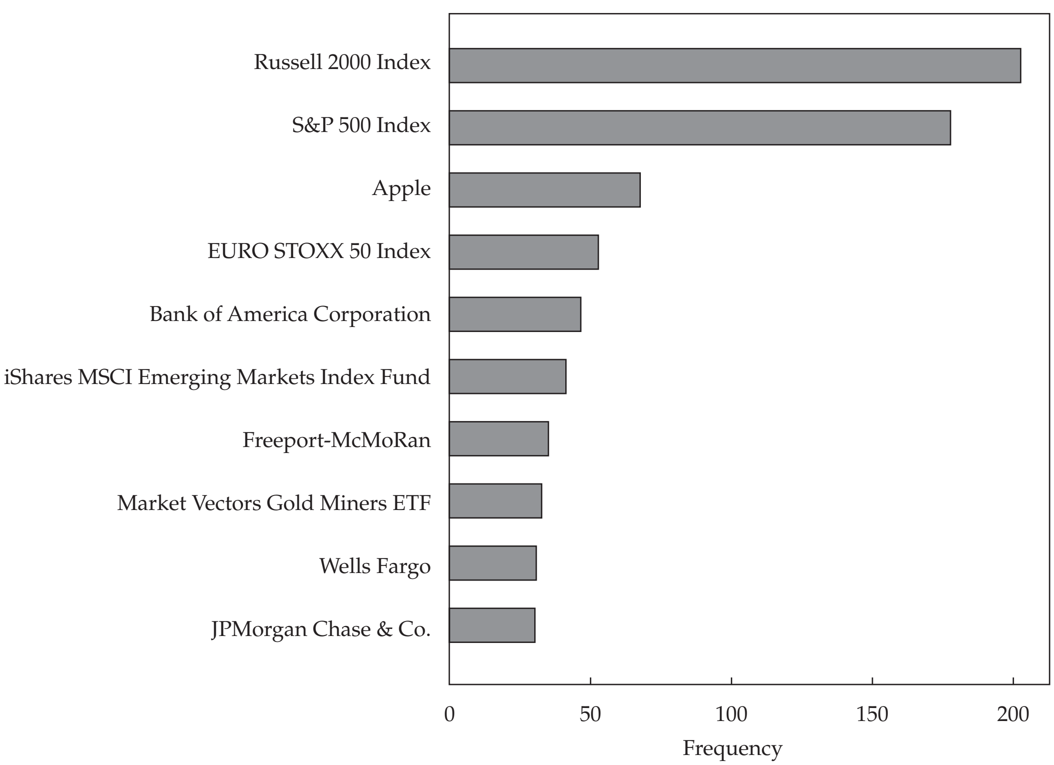

Note: The iShares Russell 2000 Index Fund is categorized together with the Russell 2000 Index.

the cases, volatility is higher than historical volatility, 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 volatility ratio.[6] This result complements Henderson and Pearson's (2011) evidence—pertaining to SPARQS—that volatility (weakly) positively affects the choice of the underlying security, but the result contradicts the evidence in Bergstresser (2008), who used a sample of all structured products and found mixed evidence on the effect of volatility on 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 high volatility, whether firm specific or marketwide. Indeed, practitioners often suggest that these securities are issued only at times of high market volatility (see Millers 2013). Later in this article, we use this evidence to justify the consideration of models with stochastic volatility to price autocalls.

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. We performed this analysis for those securities with a single underlying asset (e.g., the Russell 2000 or Apple stock). The analysis helps gauge how close the underlying securities' prices are to their most recent highs. A ratio of 1 would indicate that an offering was issued precisely at the 52-week high price. We found that 25% of the autocalls were issued within 4% of the 52-week high price and 50% were 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 repeated the analysis using a variable that adjusts for market movements: (initial price of underlying security/underlying security's 52-week high)/(S&P 500 value at issuance/S&P 500 52-week high). 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 study period supports the view that the number of autocalls issued increases during periods of market strength.[8] Similarly, Henderson and Pearson (2011) found that assets with strong performance over the past 12 months had a greater likelihood of being chosen as

Figure 3. Ratio of (Underlying Asset's Implied Volatility/Underlying Asset's Historical Implied Volatility) to (S&P 500 Implied Volatility/S&P 500 Historical Implied Volatility)

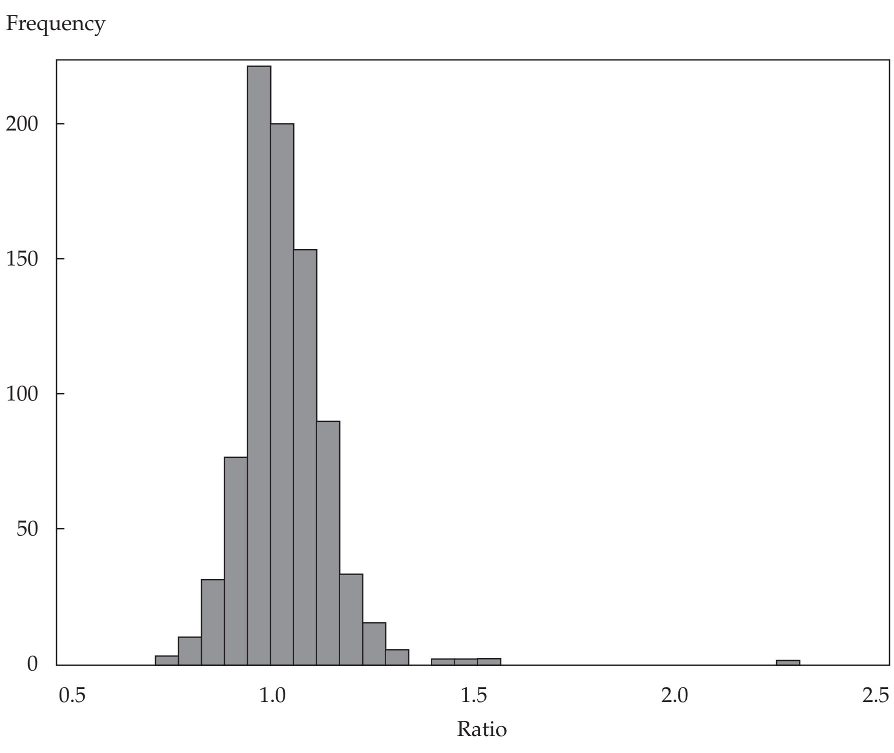

Note: Excludes autocalls with multiple underlying assets, those with the S&P 500 as an underlying asset, and those for which the implied volatility was unavailable.

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

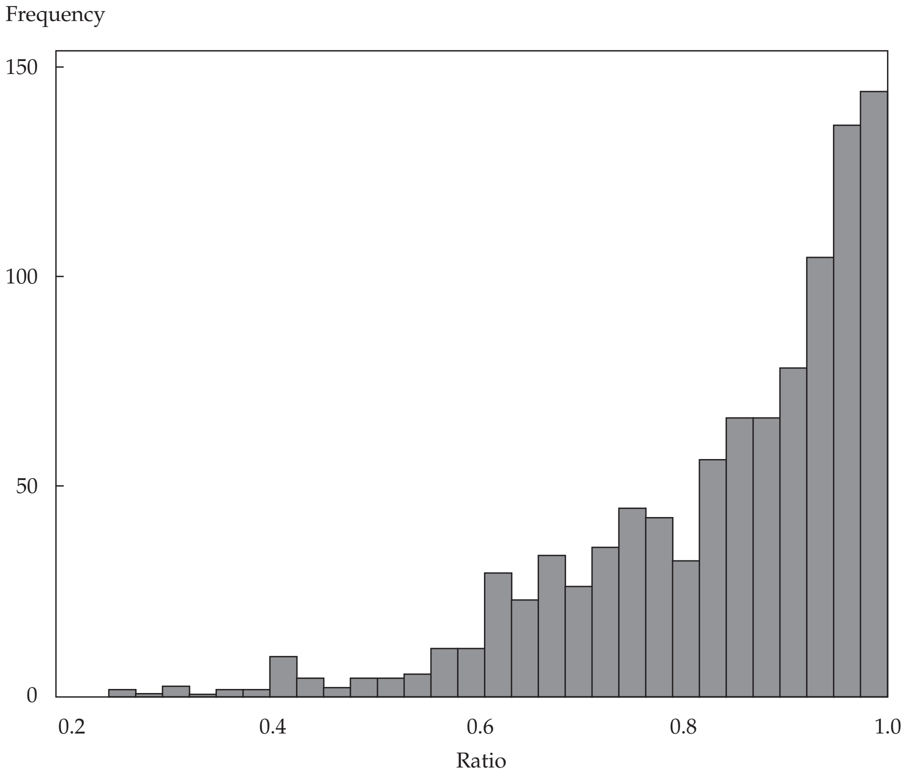

Note: Excludes autocalls with multiple underlying assets.

underlying assets. This evidence led us to consider models of mean reversion in prices when pricing autocalls. Bergstresser (2008) found mixed evidence regarding the past performance of underlying assets at issuance. We present a possible explanation for this discrepancy in the section titled "Price-Level Effects on Autocalls."

Finally, Panel J shows that mean skewness is negative (–0.31) and that 57.5% of all underlying assets display negative skewness. This finding is striking because firm-level stock returns are overwhelmingly positively skewed (see Albuquerque 2012).

Framework for Analysis

The contingent income autocallable security that we analyzed is of the plain vanilla kind. Payouts are a function of the price performance of an underlying asset as follows (Appendix A 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 is paid the coupon and principal and no further cash flows; otherwise, the security is not called, in which case the investor receives the coupon if the underlying asset's price at that date is higher than or equal to the threshold level but the investor receives nothing if the price is below the threshold. The threshold level is defined as a fraction—say, 80%—of the price at issuance. If the autocall is still alive at maturity, 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 suffers a capital loss.

Without the autocall feature, this structured product is best described as a combination of (1) 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 and (2) several short positions. These short positions include several European digital options that each mature at a different determination date and one European digital option and one European put option whose maturities coincide with that 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 determination date. It is still true, however, 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.

Geometric Brownian Motion Model of Underlying Asset Prices

In this section, we present results that describe 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. In the presence of the autocall feature, however, there is no known, exact closed-form expression for the value of the autocall security; therefore, we must resort to numerical methods to determine the properties of the investment return. To better understand the properties of the return, we simulated the properties of the plain vanilla autocall using the Apple autocall described earlier.

We calibrated several model parameters:

- The maturity of the contract is set to three 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% per year, which equals the average annualized nominal three-month T-bill rate over the period 2005–2012.
- The mean growth of the price process is set to 6.3% per year, which can be decomposed as the sum of a risk premium of 4.5% and a risk-free rate of 1.8%.
- The volatility of the price process is set to 30%, which is the implied volatility of Apple's stock return on options with one-year maturity at the time of the issue.[9]

The investor's required rate of return on the autocall is set to 6.12%. This rate equals the risk-free rate of 1.8% per year plus the risk premium of 4.5% per year times Apple's CAPM beta, which at the time of the issue was 0.96 relative to the S&P 500. The investor's required rate of return is used only to compute the fair 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 simulated 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% per year, below the assumed required rate of return of 6.12%.

Survival Probabilities. **Table 2** presents the simulated unconditional and conditional probabilities of survival. The rows labeled "baseline case" refer to simulations that use the parameters defined earlier. 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% and 25% per year). These alternative models were constructed on the basis of the evidence presented in the section titled "Autocall Sample Characteristics."

Consider the unconditional probability that the security will be called at the first determination date, which occurs if the stock price at that date is above the initial stock price. In the baseline case shown in Table 2, this probability is 51.2%. The probability that the security will be called at either the first or the second determination date 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. Because the mean of arithmetic returns is positive, the probability that the stock price will be higher than the initial stock price at the first determination date is greater than 50%. Therefore, investors who believe that the price will grow will have an incentive to invest in these securities given the high yield paid. The probability that the security will be called at the second determination date is 12.9% (equal to the probability that it will not be called by the first determination date, 1 – 0.512, times the conditional probability that it will be 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, note that the unconditional probability that it will be called at any determination date after the first year is quite low—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 in which stock price volatility is higher—equal to 40%—and all else remains equal. The increase in volatility lowers the mean of arithmetic returns to –0.017, which lowers the probability that the security will be called at the first determination date to 49.2%. In general, higher stock price volatility lowers the probability that the security will be called at any determination date and raises the probability that the security will reach maturity. The fact that the security is more likely to reach maturity implies that the stock price is more likely to be below the threshold level at maturity and hence that the investor will take a capital loss.

Table 2. Properties of Probabilities of Security Being Autocalled

Determination Date	1	2	3	4	5	6	7	8	9	10	11	12[a]
Unconditional probabilities												
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% per year)	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% per year)	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												
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% per year)	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% per year)	0.5118	0.2635	0.1778	0.1344	0.1055	0.0911	0.0789	0.0666	0.0615	0.0584	0.0466	0.0463

Notes: This table shows the model simulation of the probability that the security will be called at any determination date. The baseline model is the Apple security presented in the text. Parameters in the baseline case are as follows: risk-free rate of 1.8% per year, volatility of stock return of 30%, maturity of security of three years with quarterly determination dates, threshold level of 75%, and coupon rate of 3.525% paid quarterly.

[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 date 12 refers to maturity of the autocall.

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

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

Because of the autocall feature, a security that is called at the first determination date has a maturity that is 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 assumed 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 earlier. Panel A is common; it displays the results for the baseline case. Panels B and C depict the low- and high-volatility scenarios. Panels D and E depict the low- and high-threshold-level scenarios. Panels F and G depict the short- and long-maturity scenarios. Panels H and I depict the low- and high-coupon scenarios.

The first graph in each panel presents the unconditional frequency distribution (across all simulated paths), or histogram, of IRRs. In each histogram in every panel of Figure 5, the tallest bar represents the probability that the security will be called at the first determination date, pays the coupon, and then earns the risk-free rate up to maturity. To the right of the tallest bar, the histograms depict the events in which the security is 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 histograms. The distribution of *ex post* IRRs is considerably skewed to the left, with nonnegligible 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), when the threshold level is high (negative IRR in 12.5% of paths and IRR below –5% in 11.4% of paths), when the maturity is short (negative IRR in 11.8% of paths and IRR below –5% in 11.8% of paths), and when the coupon is low (negative IRR in 10.9% of paths and IRR below –5% in 10.6% of paths).

To better understand the negative skewness of the IRR distribution, consider the graphs on the right side of Figure 5. These plots present the average IRR across sample paths conditioned on the security surviving a given determination date. The left-most bar in each graph, indicated by determination date 0, gives the unconditional mean of the *ex post* IRR across all paths. In the baseline case, the unconditional mean IRR is 1.91% per year.

There is a general downward pattern for 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 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 a mean IRR of approximately –10% per year for those autocalls not called by the 11th determination date (33 months after issuance).

If the stock price has higher volatility, keeping all else constant, the probability that the security will be called at the first determination date decreases (see Table 2). Therefore, the likelihood that the investor will be paid the coupon decreases. Also, from Table 2, the probability that the security will survive until maturity increases. Because higher volatility in stock prices increases the probability of prices being below the threshold level, the likelihood that the investor will take a capital loss at maturity increases. IRRs decrease relative to the baseline case.[11] The effects of the increased likelihood of reaching maturity and of the price of the underlying being below the threshold level significantly reduce the mean IRR to approximately –20% per year at the 11th determination date.

Panels D and E of Figure 5 depict the effects of changing the threshold level. Consider Panel E first. The high threshold level does not affect the probability of the security being called (see Table 2) but does strictly lower the cash flows from the autocall because there are price paths that would generate

Figure 5. Simulated IRRs of the Contingent Income Autocallable Security: Baseline, Low- and High-Volatility, Low- and High-Threshold-Level, Short- and Long-Maturity, and Low- and High-Coupon Models

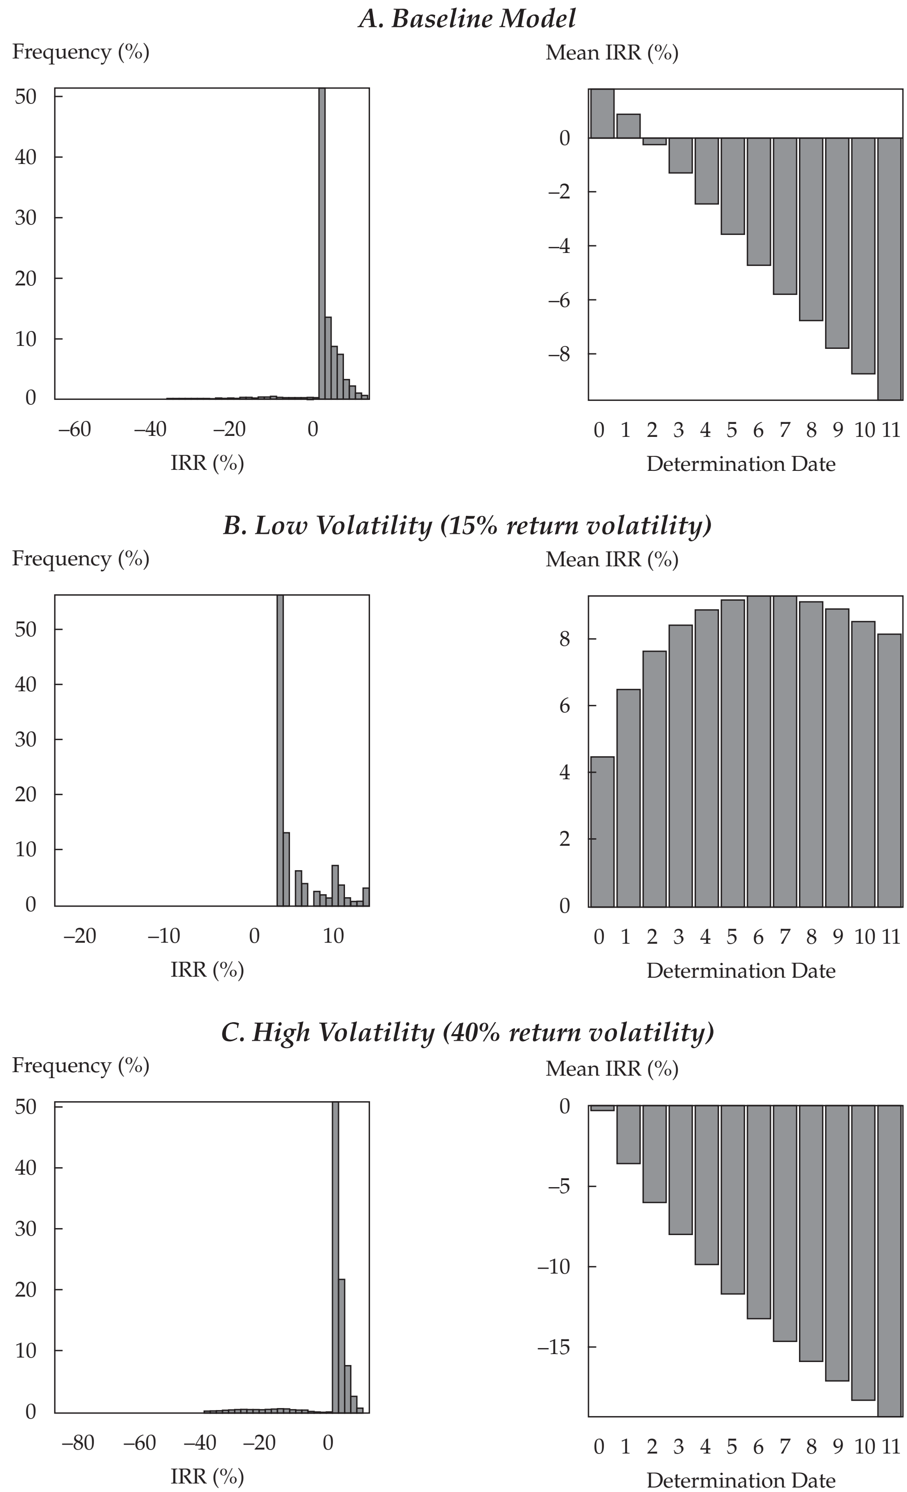

(continued)

Figure 5. Simulated IRRs of the Contingent Income Autocallable Security: Baseline, Low- and High-Volatility, Low- and High-Threshold-Level, Short- and Long-Maturity, and Low- and High-Coupon Models (continued)

D. Low Threshold Level

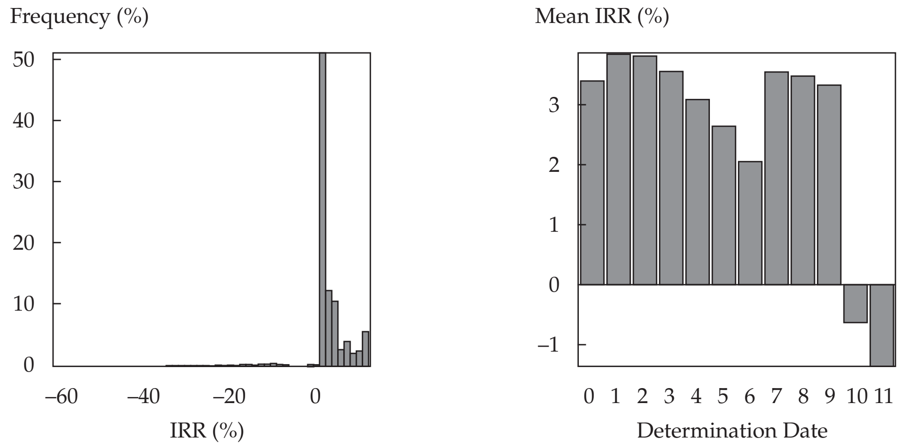

E. High Threshold Level

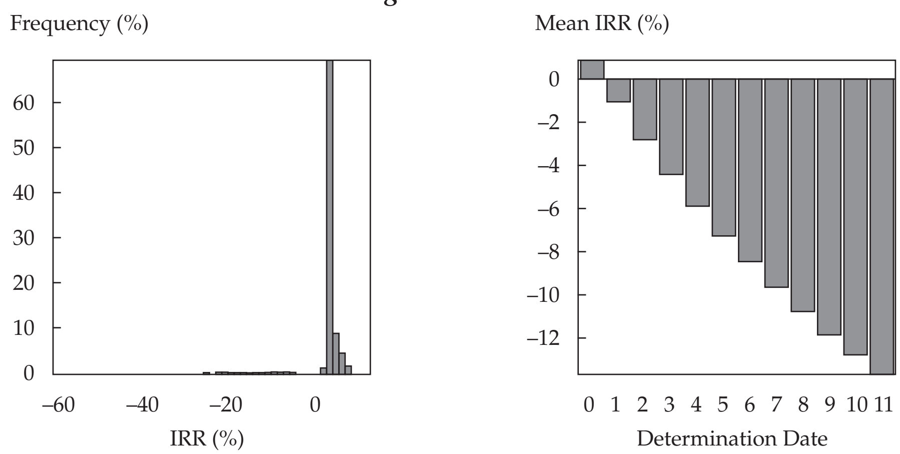

F. Short Maturity

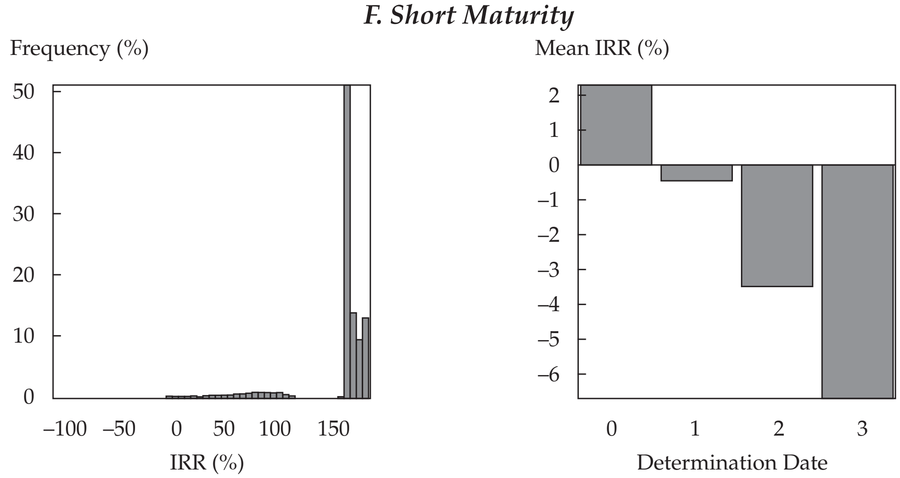

(continued)

Figure 5. Simulated IRRs of the Contingent Income Autocallable Security: Baseline, Low- and High-Volatility, Low- and High-Threshold-Level, Short- and Long-Maturity, and Low- and High-Coupon Models (continued)

G. Long Maturity

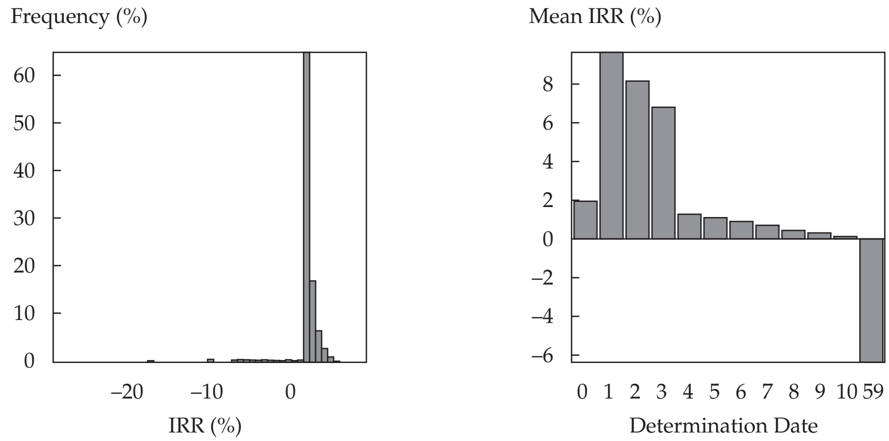

H. Low Coupon

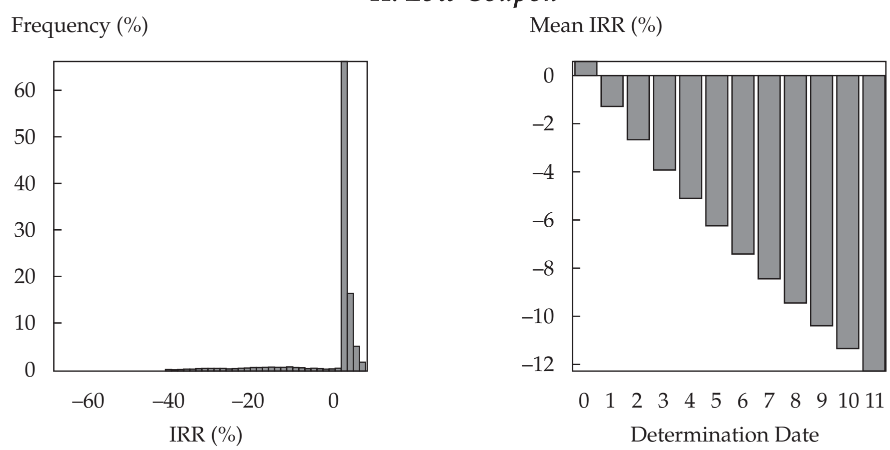

I. High Coupon

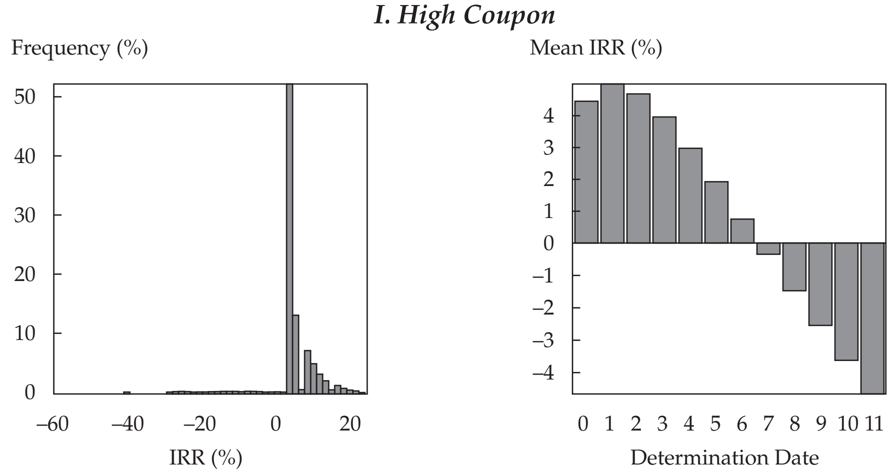

Notes: The plots are organized in the following manner. The left side of each panel contains a histogram that depicts the frequency distribution (in percent) of the annualized IRR across all simulated paths. The right side of each panel depicts the average annualized IRR (in percent) computed across all simulated paths that survive a given determination date. In the graph on the right side of Panel G, the determination dates jump from 10 to 59, the last determination date, to make the data easier to visualize.

payouts at certain determination dates that no longer generate those payouts. This is the case 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 being constant. In contrast, in the low-threshold case (Panel D), the fact that more price paths now involve a coupon payment partly offsets the negative effect of the capital loss at maturity and generates an increasing pattern of IRRs for the first few determination dates.

Panels F and G of Figure 5 depict the effects of changing maturity. Consider first Panel F. The shorter maturity does not affect the probability that the security will be called but does increase the likelihood of it reaching maturity (see Table 2). Although the probability that the stock price will be 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 –6% per year, compared 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 G) 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 (Panel H) 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 an investor who 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 will receive the underlying stock and incur a capital loss. The return in this case can be easily computed as the IRR of the following present value:

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

where S_T is the stock price at maturity T. The solution is approximately equal to IRR = $(1/12)\log(S_T/10)$. 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. This skewness reflects the investor's capital loss if the security reaches maturity that is embedded in the short option positions, and this skewness is greatly affected by several model parameters, including the volatility of the underlying asset. Investors, therefore, have much to gain from the autocall that is called early in its life and much to lose from the autocall that survives to maturity.

Volatility Exposure through Autocalls

An important assumption of the model in the previous section 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-value autocalls if this volatility is expected to decrease over time.

To operationalize these ideas, we simulated the stochastic volatility model developed by Heston (1993; Appendix A 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 includes 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 the section "Framework for Analysis" 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 practical purposes behave similarly to put options. As described previously, 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 with a low-enough threshold.

Consider a period of high volatility in which short-dated variance is greater than long-dated variance, as suggested by the evidence in the section on sample characteristics. The value of the options embedded in the autocall is lower relative to the model with constant volatility if volatility is 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 the previous section. Later in this section, we quantitatively show that these effects appear to be important determinants in the pricing of autocalls.

Heston (1993) showed that the volatility-of-volatility parameter controls the kurtosis of stock returns. Increasing the volatility of volatility has the effect of increasing kurtosis, thus generating fatter tails in stock returns, causing far-out-of-the-money put option prices to increase and near-the-money put option prices to decrease relative to the constant-volatility model. Provided the threshold level of the autocall is sufficiently high and 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 being equal.

Heston (1993) showed 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 showed that the prices of put options that are currently out of the money 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 being equal.

We quantified the significance of these effects on the pricing of the Apple autocall that we have been examining. We let all parameters common to the model in the previous section take on the same values, and therefore, we assumed that the short-dated volatility is 30%. In line with the evidence in the section on sample characteristics, we considered two values for unconditional volatility: 20% in the base-case scenario and 30%. We calibrated the mean-reversion parameter to match the mean value of persistence in asset volatility estimated by Engle and Siriwardane (2014). Following these authors, we set the mean-reversion parameter to 0.4. Besides this base case, we also report results using a higher value of mean reversion: $\lambda = 3$. When volatility reverts more rapidly to its mean, it displays less persistence. The magnitude of this parameter becomes critical when short-dated volatility and long-dated volatility differ significantly.

For 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, because we found that increasing the volatility of volatility increases the value of the autocall, we calibrated 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 Drăgulescu and Yakovenko 2002).

Table 3 presents the fair market value of the Apple autocall based on 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%. After adding a leverage effect by setting the correlation to −0.2, we found that the fair market price drops to \$9.95, representing an overpricing of 50 bps. With the largest leverage effect, a 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.

To understand the sources of gains for investors in this model relative to the model with constant volatility, we varied 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 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, which indicates that the fat tails in the model produce a positive but small effect on the price of the autocall.

Table 3. Autocall Prices in the Stochastic Volatility Model

	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

Note: Unless otherwise noted, the following are the model parameters: initial volatility of stock price of 30%; long-term 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%.

To conclude, we discuss the patterns in conditional and unconditional probabilities that the security will be 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 as well as in the baseline case (from Table 2) and in the mean-reversion case, which we will discuss in the next section. There is a decline in the probability that the security will be called after the first determination date owing to the lower long-dated volatility, but this effect is not very large given the high persistence in volatility.

Figure 6 depicts the mean IRR conditional on survival for three models: the geometric Brownian motion model of the previous section (Panel A), repeated from Panel A of Figure 5; the stochastic volatility model (Panel B); and the model of mean reversion in price from the next section of this article (Panel C). For the stochastic volatility model, the unconditional mean IRR at the time of issuance is 2.5%, higher than that in the constant-volatility model or the baseline case. This finding reflects the higher valuation and lower overpricing discussed earlier. The mean IRR conditional on survival declines as the expected capital loss at maturity increases. In the paths where the security is never called (the right-most bars), the mean IRR is negative but less so than for the baseline model. This finding reflects the lower value of the options (including the put option at maturity) due to the mean reversion of volatility, which 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.

Price-Level Effects on Autocalls

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). (Appendix A 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 calibrated the mean-reversion parameter to 0.75 and the long-term mean of the price to 20% below the price at the issuance date. According to this long-term mean, the seller of the autocall believes that there is a significant probability that the stock price will fall. In addition, considering 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 motion model, we adjusted 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% per year—significantly lower than the required rate of 6.12%. The reason that the discount is so much larger relative to that in the baseline model of the "Geometric Brownian Motion Model of Underlying Asset Prices" section of this article (1.4%) and that the IRR is lower is that the stock price at the issuance date is significantly higher than its long-term mean of \$8. This fact has two main consequences: (1) the security is less likely to be called and to pay coupon plus interest at the first determination date, and (2) although it pays interest whenever the stock price is between \$7.50 and \$10, it is more likely to drop below the threshold and result in a capital loss for the investor (note that the long-term mean of the stock price is \$8 and the threshold is \$7.50). As shown in Table 4, the unconditional probability of the security being called at the first determination date is only 37.2% and at either 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 will reach maturity is 29% (up from 15% in the baseline case).

Table 4. Properties of Probabilities of Security Being Autocalled

Determination Date	1	2	3	4	5	6	7	8	9	10	11	12[a]
A. Unconditional probabilities												
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
B. Conditional probabilities												
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	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

Notes: This table shows the model simulation of the probability that the security will be called at any determination date. The baseline model parameters are as indicated in Table 2. The following are the stochastic volatility model parameters: initial volatility of stock price of 30%, long-term 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%. The mean-reversion model parameters are volatility of stock price of 31.7%, long-term mean price 20% below the price of the underlying at issuance, and parameter controlling the speed of adjustment of 0.75.

[a]Refers to the probability of reaching maturity.

Figure 6. Simulated IRRs Based on Different Underlying Asset Pricing Models

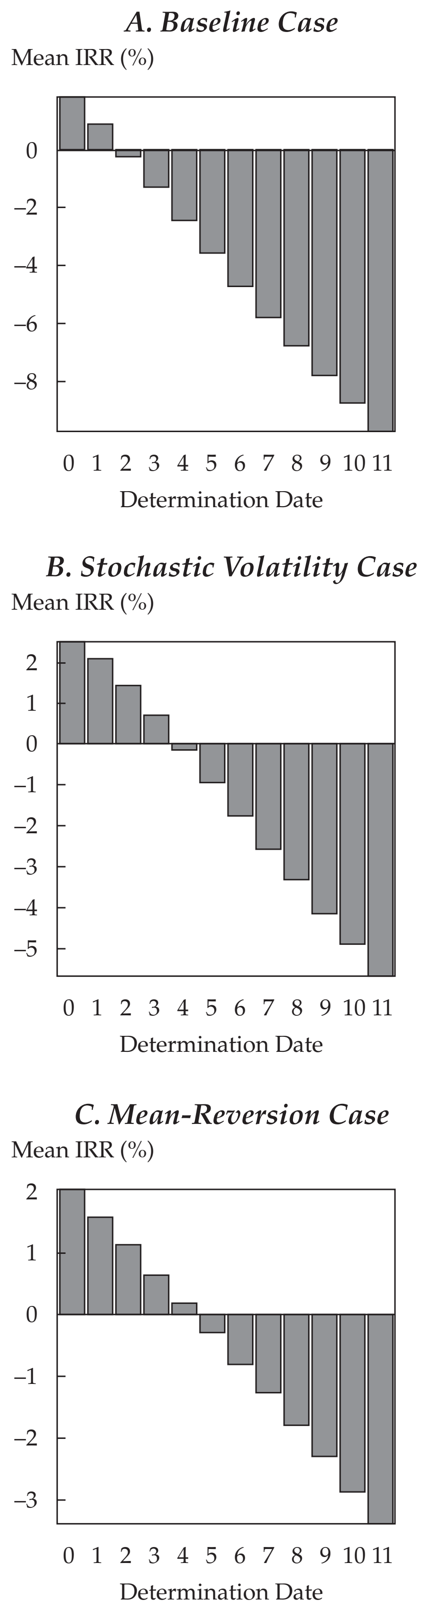

Note: This figure shows mean annualized IRRs by determination date in percent.

In Panel C of Figure 6, we plot the mean IRR conditional on survival. The unconditional mean IRR at the time of issuance (left-most bar) is 2.1% in the mean-reversion model, higher than the baseline model's 1.91%. As 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 to the past performance of the underlying asset should depend on whether the structured security has embedded call- or put-like options. Because Bergstresser pooled both call- 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.

Conclusion

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

We offer two main takeaways. First, we documented 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 generally performs well in the stock market, displaying prices at or near the 52-week high value. Second, we used this evidence and evidence from other financial markets to show that the most common model used in the literature on valuing structured products, geometric Brownian motion, is inappropriate for this task. We have shown 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.

Therefore, it is possible that investors' views on the underlying assets regarding price and volatility dynamics may justify the overall interest in these securities and the growth in the market. Notwithstanding the finding that an appropriate pricing model can approximate the fair value of the autocall that is close to its actual price, private wealth managers and financial advisers 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 Brothers case suggests, it can be large.

We conclude with one final remark to financial analysts. Issuing banks engage in significant hedging of the exposures created by selling these structured products. Bennett and Gil (2012) cautioned of a potential "vicious circle" in which a price decrease 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 derivative markets that can be used for hedging or at times when liquidity in these markets dries up (see Millers 2013).

We thank Kyle Darres for his excellent research assistance, Rachael Christensen and Scott McEwan of Morgan Stanley for providing data, and Harvey Boshart, CFA, of 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. In addition, Rui Albuquerque acknowledges the Portuguese Science Foundation for financial support under the project PTDC/EGE-GES/120282/2010, and research conducted by Raquel Gaspar was partially supported by the Portuguese Science Foundation under the SANAF project UTA_CMU/MAT/0006/2009.

CE Qualified Activity CFA Institute: 1 CE credit.

Appendix A. Framework for Analysis

The contingent income autocallable security that we analyzed 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 issuance date is S_0. For simplicity, the price and notional value of the autocall is $P = S_0$.
- At determination date $t = 1, \ldots, T-1$, the security is called if $S_t > S_0$, in which case the investor gets $(1 + i)P$, where i is the coupon rate, and no further cash flows; otherwise, it is not called, in which case the investor gets iP if $S_t > \alpha S_0$, with $\alpha < 1$ (αS_0 is the threshold level).
- At maturity, the autocall pays $(1 + i)P$ if $S_T > \alpha S_0$; 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 studied the 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:

$$\frac{S_t}{S_0} = \exp\left[\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma W_t\right], \tag{A1}$$

where

μ = the instantaneous growth rate in prices

σ = the instantaneous return volatility

W_t = a Wiener process whose continuous increments are normally distributed with a zero mean and unit variance

To simulate the process, we discretized the process in Equation 1 using[13]

$$\frac{S_{t+\Delta t}}{S_t} = \exp\left[\left(\mu - \frac{1}{2}\sigma^2\right)\Delta t + \sigma N(0,1)\sqrt{\Delta t}\right]. \tag{A2}$$

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

$$\frac{S_{t+\Delta t}}{S_t} = \exp\left[\left(\mu - \frac{1}{2}v_t\right)\Delta t + W_t\sqrt{v_t \Delta t}\right],$$

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

$$v_{t+\Delta t} = v_t + \lambda\left(\bar{v} - v_t\right)\Delta t + \theta Z_t \sqrt{v_t \Delta t} + \frac{1}{4}\theta^2 \Delta t\left(Z_t^2 - 1\right). \quad \text{(A3)}$$

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, W_t and Z_t, respectively, are assumed to be correlated, with a correlation coefficient of ρ; v_t is the conditional variance of stock returns, which we call "short-dated variance"; and $\bar{v}$ is the long-term mean of the conditional variance, which we call "long-dated variance."

The third and final model we studied 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):

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

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-term mean of the price.

Notes

1. In 2012, more than 450 Apple-linked structured products were brought to market, with at least 75% of them issued when the stock's price was at least \$550 (Zweig 2013). We found 68 autocalls issued on Apple through mid-2013, based on US SEC security registration data.
2. The preliminary prospectus indicated 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 www.sec.gov/Archives/edgar/data/895421/000095010312003734/dp31778_fwp-ps257.htm.
3. Using a dataset from Bloomberg that includes autocalls issued outside the United States, Deng, Mallett, and McCann (2011) found that the autocall market 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 United Kingdom are often based on the FTSE 100 Index, whereas those in continental Europe are often based on the EURO STOXX 50 Index or other major indexes from Europe and elsewhere. Similar to autocalls in the United States, 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, Uridashi autocalls have been popular instruments for retail investors 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 benefited from the low interest rates available to Japanese yen deposits and from the fact that they are denominated in a foreign currency.
4. Szymanowska, Ter Horst, and Veld (2009) used a constant elasticity variance model, but as in the geometric Brownian motion model, volatility is constant in their model.
5. The SEC search permits only a four-year historical window to be examined. Using the Bloomberg database, we were able to find 29 autocalls registered with the SEC prior to 2009, representing \$226 million. Because prior to 2009 there were so few autocalls, we constrained our data to the autocalls that have Form 424(b)(2) available in EDGAR to ensure the comparability and quality of the data. Deng et al. (2011) instead searched the Bloomberg database for autocalls and were able to obtain a significantly larger sample than ours. Their larger sample size can be explained mostly by the inclusion of autocallable securities issued outside the United States.
6. The general tone of these results was unchanged when we used realized historical return data and robust measures of variance—the interquartile range and the median absolute deviation. Results are available from the authors upon request.
7. Available from the authors upon request.
8. Oyo (2013) suggested that the improved prospects of the US economy following the crisis drove the use of US indexes as underlying assets, whereas the deteriorating prospects of European economies drove usage away from the more stagnant European-based indexes (for information on market shares, see Thin 2014).
9. To compute implied volatility, we chose one year for the option maturity, instead of three years, to match the maturity of the autocall owing to a lack of liquidity on long-dated options.
10. Alternatively, we calculated 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 this article 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 space, we omitted the presentation of the results under this alternative approach to calculating the *ex post* IRR.
11. In the low-volatility case, the mean IRRs display a hump-shaped pattern. 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 will receive the coupon, which is better than having the security be called earlier 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 will be below the threshold will be large enough to put downward pressure on the IRRs and to give rise to the decreasing part of the hump.
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 will receive the coupon at 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.
13. Deng et al. (2014) provided an approximate analytical solution to the price of an autocall.

14. The discretized variance equation (Equation A3) in the Heston model does not guarantee that variance is nonnegative. Several approaches have been proposed to minimize this concern. In this article, we used the Milstein scheme (see Kahl and Jäckel 2006). The Milstein scheme adds the last term on the right-hand side of the variance equation. We simulated the model with various discretizations, with similar results across approaches.

References

Albuquerque, Rui. 2012. "Skewness in Stock Returns: Reconciling the Evidence on Firm versus Aggregate Stock Returns." *Review of Financial Studies*, vol. 25, no. 5 (May): 1630–1673.

Bennett, Colin, and Miguel A. Gil. 2012. "Volatility Trading." Working paper, Grupo Santander.

Bergstresser, Daniel. 2008. "The Retail Market for Structured Notes: Issuance Patterns and Performance, 1995–2008." Working paper.

Bernard, Carole, Phelim P. Boyle, and William Gornall. 2011. "Locally Capped Investment Products and the Retail Investor." *Journal of Derivatives*, vol. 18, no. 4 (Summer): 72–88.

Breuer, Wolfgang, and Achim Perst. 2007. "Retail Banking and Behavioral Financial Engineering: The Case of Structured Products." *Journal of Banking & Finance*, vol. 31, no. 3 (March): 827–844.

Burth, S., T. Kraus, and H. Wohlwend. 2001. "The Pricing of Structured Products in the Swiss Market." *Journal of Derivatives*, vol. 9, no. 2 (Winter): 30–40.

Célérier, Claire, and Boris Vallée. 2014. "What Drives Financial Complexity? A Look into the Retail Market for Structured Products." Working paper.

Chang, Eric C., Dragon Yongjun Tang, and Miao Zhang. Forthcoming. "Suitability Checks and Household Investments in Structured Products." *Journal of Financial and Quantitative Analysis*.

Coval, Joshua, Jakub Jurek, and Erik Stafford. 2009. "The Economics of Structured Finance." *Journal of Economic Perspectives*, vol. 23, no. 1 (Winter): 3–25.

Das, Sanjiv, Seoyoung Kim, and Meir Statman. 2013. "Liability Directed Investing in a Behavioral Portfolio Theory Framework." Working paper.

Deng, Geng, Tim Dulaney, Tim Husson, Craig McCann, and Mike Yan. 2014. "Ex Post Structured Product Returns: Index Methodology and Analysis." Working paper.

Deng, Geng, Guo Huali, and Craig McCann. 2009. "Structured Products in the Aftermath of Lehman Brothers." Working paper.

Deng, Geng, Joshua Mallett, and Craig McCann. 2011. "Modeling Autocallable Structured Products." *Journal of Derivatives & Hedge Funds*, vol. 17, no. 4 (November): 326–340.

Dixit, Avinash K., and Robert S. Pindyck. 1994. *Investment under Uncertainty*. Princeton, NJ: Princeton University Press.

Drăgulescu, A.A., and V.M. Yakovenko. 2002. "Probability Distribution of Returns in the Heston Model with Stochastic Volatility." *Quantitative Finance*, vol. 2, no. 6: 443–453.

Engle, Robert F., and Emil N. Siriwardane. 2014. "Structural GARCH: The Volatility–Leverage Connection." Working paper.

Fabozzi, Frank J. 2005. "The Structured Finance Market: An Investor's Perspective." *Financial Analysts Journal*, vol. 61, no. 3 (May/June): 27–40.

Financial Conduct Authority. 2013. "Structured Products" (www.fca.org.uk/firms/financial-services-products/investments/structured-products).

Henderson, Brian J., and Neil D. Pearson. 2011. "The Dark Side of Financial Innovation: A Case Study of the Pricing of a Retail Financial Product." *Journal of Financial Economics*, vol. 100, no. 2 (May): 227–247.

Hens, Thorsten, and M. Rieger. 2008. "The Dark Side of the Moon: Structured Products from the Customer's Perspective." Working paper.

Heston, Steven L. 1993. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *Review of Financial Studies*, vol. 6, no. 2 (April): 327–343.

Kahl, Christian, and Peter Jäckel. 2006. "Fast Strong Approximation Monte-Carlo Schemes for Stochastic Volatility Models." *Quantitative Finance*, vol. 6, no. 6: 513–536.

London Stock Exchange. 2014. "Structured Products" (www.londonstockexchange.com/specialist-issuers/structuredproducts/sp-introbrochure.pdf).

Millers, Vita. 2013. "Korea Autocalls Pose Risk to Structured Products Market." Risk.net (17 September).

Oyo, Hiroshi. 2013. "SocGen: Japanese Uridashi Market Growing Despite Challenges." Structured Retail Products (15 October): www.structuredretailproducts.com/blog-story/17596/socgen-japanese-uridashi-market-growing-despite-challenges#.VHD-xtapq0M.

Pereira da Silva, Paulo, and Fernando Silva. 2013. "The Hidden Costs of Structured Retail Products—Evidence for the Portuguese Market." Working paper.

Purnell, Freya. 2012. "Why Financial Advisers Are Avoiding Structured Products." *Money Management* (8 November): www.moneymanagement.com.au/news/financial-planning/why-financial-advisers-are-avoiding-structured-products.

Stein, J. 2013. "Overheating in Credit Markets: Origins, Measurement, and Policy Responses." Unpublished speech, Federal Reserve Bank of St. Louis.

Stoimenov, Pavel A., and Sascha Wilkens. 2005. "Are Structured Products 'Fairly' Priced? An Analysis of the German Market for Equity-Linked Instruments." *Journal of Banking & Finance*, vol. 29, no. 12 (December): 2971–2993.

Szymanowska, Marta, Jenke Ter Horst, and Chris Veld. 2009. "Reverse Convertible Bonds Analyzed." *Journal of Futures Markets*, vol. 29, no. 10 (October): 895–919.

Thin, Win. 2014. "The Carry Trade, Emerging Markets and Japanese Retail Investors." *Credit Writedowns* (5 February): www.creditwritedowns.com/2014/02/carry-trade-emerging-markets-japanese-retail-investors.html.

Wilkens, S., C. Erner, and K. Roder. 2003. "The Pricing of Structured Products in Germany." *Journal of Derivatives*, vol. 11, no. 1 (Fall): 55–69.

Zweig, Jason. 2013. "How Apple's Fall Bit Bondholders, Too." *Wall Street Journal* (26 January).