Locally-Capped Investment Products and the Retail Investor

Carole Bernard
University of Waterloo
c3bernar@uwaterloo.ca

Phelim Boyle*
Wilfrid Laurier University
pboyle@wlu.ca

Abstract

Structured products are popular with retail investors. Many of them provide a guaranteed return combined with some participation in the performance of the equity market. These contracts often have complex designs. We focus on a particular design where the investor's return is capped periodically and there is also a guaranteed minimum at maturity. Standard finance theory predicts that consumers should prefer simpler contracts. If consumers overweight the probability of getting the maximum possible return they may prefer the more complex contract and we provide evidence that sellers encourage this type of overweighting. We also discuss other explanations for this puzzle.

Keywords: Household finance, structured products, retail investor, locally-capped contracts.

*We thank Will Gornall for excellent research assistance. We would also like to thank Peter Carr, Daniel Dorn, Philip Dybvig, Richard Green, Mary Kelly and Harry Panjer for helpful discussions. Thanks also to participants at the following conferences: the Northern Finance Association Annual Meeting in Kananaskis Village, Canada, the Euro-Working Group in Financial Modeling at the Cass Business School in London. We are grateful for comments from seminar participants at the University of Illinois at Urbana-Champaign, Wilfrid Laurier University, the Chinese Academy of Sciences in Beijing, Leeds University Business School, Lyon Business School and Cass Business School. Both authors acknowledge support from the Natural Sciences and Engineering Research Council of Canada. C. Bernard also acknowledges support from Tata Consultancy Services.

1 Introduction

The field of household finance is concerned with the study of how households make investment decisions to attain their objectives. Campbell (2006) in his presidential address highlighted the importance of studying this neglected field. There is often considerable disparity between the predictions of the standard finance models and what households actually do in practice. Campbell notes that the distinction between positive and normative models poses a particular challenge in the context of household finance and suggests that these discrepancies are central to the field of household finance.

The current paper discusses these issues in the context of consumer choice in the market for structured retail financial products. This is an important market and these products represent one of the ways in which many investors gain exposure to equity markets. The typical contract has a payoff based on the performance of an equity portfolio[1], so that the investor benefits from good stock market performance. At the same time, the contract often provides a minimum guaranteed floor that gives protection against poor equity returns. The cost of the guarantee is generally covered by modifying the payoff for example capping it at some level. Such products have a strong consumer appeal since they combine upside appreciation in bull markets with downside protection in bear markets in a single unified contract.

The simplest versions of such a contract can be viewed as a portfolio of a long position in a zero-coupon bond and a call option on the underlying equity portfolio. While there are some contract designs similar to this in the market, the vast majority of contracts have more complicated designs. For example, the terminal payoff may be based on the sequence of realized returns on the equity portfolio where each return is capped at a specific level. The

[1]Such as the S&P 500 for about 45% of index linked products issued from 1992 through 2005 (see Table 2 of Henderson and Pearson (2007) for more details).

floor guarantee can be based on the entire life of the contract or there could be a minimum return each year of, say, zero per cent. Other aspects of the design make these contracts difficult for the typical investor to understand. It is hard to make price comparisons across the different products.

In this paper, we describe and analyze a popular class of structured products. These contracts are subject to periodic caps with a maturity guarantee and are known as *locally-capped, globally-floored contracts*. We show that they are dominated by consistently priced contracts with a single cap at maturity and the same type of floor. Campbell (2006) and Carlin (2008) observe similar inefficiencies in other financial decisions made by consumers.

Henderson and Pearson (2007) also find puzzling investor behavior in the structured products market and note "*it is difficult to rationalize investor demand for structured equity products within any plausible normative model of the behavior of rational investors.*"

There is evidence[2] that retail structured products tend to be overpriced and also complex. The existence of overpriced complex financial products is consistent with Carlin's (2008) model. In this model, firms strategically increase the complexity of financial products to preserve market power and bound the financial literacy of consumers. Moreover sellers are able to charge a higher markup on the more complex products. In most cases, the retail investor does not have the expertise to understand the complexities of these contracts and obtains advice from an agent who is remunerated by sales commissions. If the producer's surplus is shared with the sales agent, then there are incentives for the agent to push the more complex products and for the industry to favor a regulatory regime that makes it easier to avoid disclosure and complicate the product. When fees and commissions are taken into account, structured products are very expensive. On the other hand, these instruments provide consumers with features that would be difficult for them to replicate on their own and

[2]Stoimenov and Wilkens (2005), Henderson and Pearson (2007)).

the argument is sometimes made that they contribute to completing the market (Rossetto and Van Bommel (2008)). Consumers are also attracted by prepackaged solutions (Shefrin and Statman (1993), Shefrin (2000)).

In this paper, we first show with specific examples that the simpler products should be preferred by investors in contrast to what is observed in the market. From a theoretical perspective, the result is more general: it is not optimal for an agent who maximizes expected utility of final wealth to invest in path-dependent securities. This is a consequence of the concept of cost-efficiency introduced by Cox and Leland (2000)[3] and Dybvig (1988a, 1988b). One way to rationalize these preferences is to go beyond the expected utility framework. We can explain the demand for locally-capped contracts by modifying standard consumer preferences using ideas of Tversky and Kahneman (1992). We show that the overweighting of the extreme probabilities discussed also by Barberis and Huang (2008) can be used to explain the demand for locally-capped globally-floored contracts. The maximum possible return under these contracts has a negligible probability of being attained. But if this low probability event is included as one of a few possible projected scenarios, this may influence the buyer's perception of it occurring. This overweighting can lead a consumer to value these products more highly relative to comparable contracts. Of course there may be other reasons that influence the demand for locally-capped design.

The layout of the rest of this paper is as follows. In the next section, we describe some general features of the structured products market and some specific types of contracts. Section Three focuses on locally-capped and globally-capped contracts with a maturity guarantee. We show that the decision to invest in these contracts is puzzling from the standpoint of normative investment theory. We are able to quantify the prices and risk return profiles of these contracts under plausible assumptions and compare them with simpler contracts

[3]An earlier version of this paper dates back to 1982 in the Proceedings of the seminar on the Analysis of Security Prices, Center for Research in Security Prices, University of Chicago.

that are also available to investors. Section Four discusses how this puzzle can be at least partly explained by using the overweighting approach. We show that this phenomenon is also consistent with the predictions of the Carlin (2008) model. Section Five provides evidence of how investors could be induced to overweight the probability of getting the maximum possible return. Section Six contains a short summary.

2 The Structured Products Market

In this section, we provide a brief overview of the structured products market in the United States. Then, we describe the locally-capped globally-floored contracts that we are interested in. We illustrate the features of these contracts using actual examples.

There are different ways to classify structured products. We can classify them by the type of seller, the way in which they are regulated and of course the characteristics of the products. There are two main groups of providers of structured investment products to the retail investor: investment banks and insurance companies. Contracts sold by banks are listed on exchanges and are pure investment products. Exchange listed products are regulated by the SEC. The structured products sold by insurance companies are packaged with other features such as insurance benefits. They include Variable Annuities (VA) and Equity Indexed Annuities (EIA). Equity Indexed Annuities are regulated by state insurance commissioners. On the other hand, Variable Annuities are viewed as hybrid products with both insurance and investment features and are regulated both by the SEC as securities and by the various state insurance commissions as insurance contracts.

The size of the structured products market has expanded significantly in recent years. Henderson and Pearson (2007) note that the dollar value of new exchange listed products grew from about one billion dollars in 1997 to about ten billion in 2005. Equity Indexed

Annuities (EIAs) sold by insurance companies have grown more dramatically. EIAs were first introduced in 1995 and sales in 2006 were about 25 billion dollars (Palmer (2006)). The variable annuity market is much larger than these two: it is estimated [4] that as of December 2006 the total assets under management in variable annuity accounts was about 1.4 trillion dollars. In this paper, we focus on exchange listed products sold by banks.

Most of these products are listed on the American Stock Exchange (AMEX). There are five categories of structured products: Index-Linked Notes, Equity-Linked Term Notes, Currency Notes, Index-Linked Term Notes and others. We focus on Index-Linked Notes since these products represent the largest share (at least 60%) of the volume among all structured products. As of Oct 3, 2006, there were 208 Index-Linked Notes amounting to $7.9 billion, and 320 structured equity products listed on the AMEX. We focus on the locally-capped, globally-floored design.

At this stage it is helpful to describe some of the main characteristics as well as certain institutional details of the exchange listed, locally-capped, globally-floored contracts. There are 39 such contracts listed on the AMEX and Appendix A summarizes their main features. The total dollar volume is 2.9 billion and the most commonly used reference index (14 out of 39) is the S&P 500. The average maturity[5] of these 39 contracts is about 5.3 years. The cap frequency is either monthly (11 contracts) or quarterly (28 contracts). The size of the average cap is about 4.6 percent per month for the monthly-capped contracts and about 8 percent per quarter for the quarterly-capped contracts.

[4] See Evans and Fahlenbrach (2007).

[5] Here we mean time to maturity measured from the start of each contract.

2.1 The Quarterly Sum Cap

We now describe three quarterly-capped contracts issued by J.P. Morgan Chase and listed on the AMEX. They were issued originally as five year notes maturing in 2009. Table 1 gives some summary details[6].

Table 1: Details of three J.P. Morgan's Quarterly Sum Cap Contracts

Ticker Symbol	Index	Proceeds \$ millions	Commission Rate	Cap Level	Cap frequency	Maturity Guarantee
JPL.E	S&P	21.2	5.00%	5 %	Quarterly	10%
JPL.G	S&P	22.5	6.31%	6 %	Quarterly	10%
JPL.H	S&P	4.1	3.00%	6 %	Quarterly	10%

All three tranches have identical terms to maturity (five years). In each case, the reference index is the S&P Index. Each product has the same global guarantee. If an investor invests \$1,000 she is guaranteed to get back \$1,100 at maturity irrespective of what happened to the S&P Index. The Index returns do not include dividends. The Pricing Supplements issued in connection with these securities provide an extensive discussion of how the returns are computed and in particular how the cap feature works. For example, in the case of the $JPL.G$ contract the Pricing Supplement dated June 22, 2004 describes the payoffs in the first paragraph as follows:

> *At maturity, J.P. Morgan Chase & Co. will pay in cash, for each \$1,000 principal amount of notes, the greater of \$1,100 which we refer to as the minimum payment amount, and \$1,000 plus an amount, which we refer to as the additional amount, based on the capped quarterly upside performance and full downside quarterly*

[6] More complete details are available from the Pricing Supplements issued in connection with these securities. The Pricing Supplement for the $JPL.E$ tranche is dated December 12, 2003. The Pricing Supplement for the $JPL.G$ tranche is dated June 22, 2004. The Pricing Supplement for the $JPL.H$ tranche is dated September 16, 2004. These supplements can be downloaded from the website `www.amex.com`.

performance of the S&P 500 index.

The operation of the quarterly cap is explained more fully later on the same page:

> *This 6% cap on any quarterly capped Index return means that your ability to benefit from an increase in the Index is limited to 6%, regardless of how well the Index performs in any quarterly valuation period. Conversely, there is no limit to any decreases in the quarterly capped Index return which means you have no protection against negative performance of the Index in calculating the additional mount.*

The Pricing Supplement describes how the returns are calculated.

> *The additional amount will be calculated by the calculation agent by multiplying $1,000 by the sum of the quarterly capped Index returns for each of the 20 quarterly valuation periods during the term of the notes.*

The Pricing Supplement also includes five hypothetical projections which illustrate how the payment at maturity of a Quarterly Sum Cap is calculated under different market scenarios. Details of the first two examples are given in Appendix C. The first projection assumes the investor receives the maximum possible return under the contract. Since the cap rate is 6% per quarter, the return over 5 years of the quarterly-capped contract can theoretically reach 120% if over the twenty consecutive quarters, the index has a return higher than (or equal to) the cap level 6% (120% is the sum of twenty returns of 6%). We note that the inclusion of the maximum possible return as a possible scenario may serve to influence a buyer's perception of the probabilities. We return to this issue later.

We can price these contracts using standard valuation methods. Assume x_0 is the initial investment. We assume a risk-free rate of 5% and a dividend rate of $\delta = 2\%$. We assume that the minimum guaranteed rate is $g = 10\%$ and that the maturity is $T = 5$ years. The contract can be viewed as a portfolio of a zero coupon bond with maturity amount $(1+g)x_0 = \$1,100$ and an additional option component. Figure 1 shows how the price of the option component depends on the volatility of the index for a range values.

Figure 1: Value of option component under Quarterly Sum Cap.
Prices of the option component (based on the difference between the payoff and the guarantee $(1+g)x_0$ as a function of the assumed volatility (between 5% and 40%). The cap rate is 6% and the contract matures in 5 years. The dividend rate $\delta = 2\%$, the risk-free rate is $r = 5\%$.

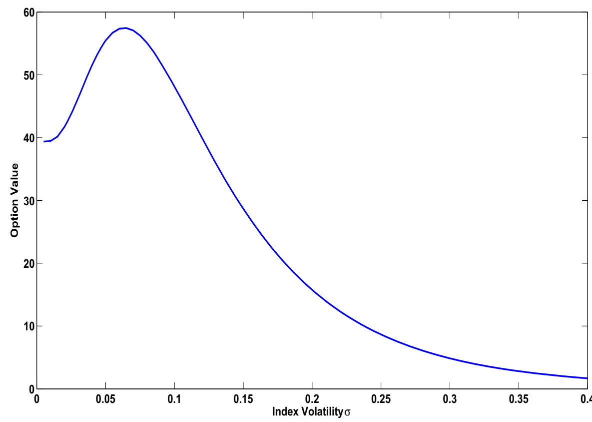

Assume an investor invests \$1,000 in the security at time zero. The no-arbitrage value of the guarantee at time zero is equal to $\$1,100\, e^{-5r} = \$856.7$ when $r = 5\%$. From Figure 1, the option component attains a maximum value slightly lower than \$60 when the volatility of the index is about 8%. As the volatility increases the value of the option component declines because of the cap feature. The commission rate for the $JPL.G$ tranche is 6.31%. In this case, the institution's minimum profit is at least $1000 - 856.7 - 60 - 63.1 = 20.2$ (when the volatility is 8%). Of course, the bank will incur other expenses in running the business. Moreover, Wilmott (2002) shows that it may not be possible to hedge the embedded cliquet option for its no-arbitrage price since this option can be expensive to hedge. The commission rate for the smaller $JPL.H$ tranche is 3.0% (see table 1) which implies a higher profit for the institution.

This result is consistent with Hendersen and Pearson (2007) who found that structured products can be significantly overpriced with an extra premium of 8% between their model price and the price charged in the market.

3 Analysis of Locally and Globally-Capped Contracts

In this section, we analyze the locally-capped globally-floored contracts. To better understand the consequences of the local cap (monthly or quarterly), we consider two contract designs: the *globally-capped contract* and the *locally-capped contract.* Both contracts have a global floor (a guarantee at maturity). For simplicity we refer to the first design as the globally-capped (GC) design and the second as the locally-capped (LC) design with the understanding that both contracts have a guarantee. The major difference is that the LC design is highly *path-dependent* whereas the GC design is a *point-to-point* design that depends only on the beginning and ending values of the underlying index.

We use two approaches to make the theoretical comparison. The first approach allows us to exhibit a simple contract that dominates the LC contract for all investors with minimal assumptions on agents' preferences. However this approach does not permit a direct comparison between the LC contract and the GC contract. The second approach assumes specific preferences and compares the two designs. Both approaches show that the LC design is not optimal for investors whose objective is to maximize the expected utility of final wealth.

The first approach applies arguments of Dybvig (1988a,1988b) on the efficiency cost of investment strategies. We compare the efficiency cost of investing in both contracts. Under fairly general assumptions (independent of individuals preferences), we show that it is never optimal for an agent who maximizes expected utility of final wealth to invest in a path-dependent contract. Specifically, we show that the LC contract is dominated by a *point-to-point* design. Furthermore, we can quantify the efficiency loss of investing in the LC contract.

In the second approach, we compare the LC design with the GC design when these contracts are both fairly priced (and have the same cost). We show that it is never optimal for agents with mean variance preferences to buy the LC contract product if the GC product is available and fairly priced[7]. In the case of more general preferences, we will show that for realistic values of agent's risk aversion, agents will prefer the globally-capped contract when the volatility is sufficiently high.

Our goal is to describe this puzzle: namely the consumers preference for the complex LC contract instead of the simpler GC contract. We assume the initial investment is $x_0 = \$1,000$. Although we use numerical examples based on one set plausible parameter values, the conclusions are robust to other parameter choices.

3.1 Description of Contracts

This section describes the essential features of the LC contract and the GC contract and introduces our notation and conventions.

Let C be the global cap at the maturity of five years ($T = 5$). Let g be the minimum guaranteed rate. Let S_t be the value of the underlying index at time t. The maturity payoff under the GC design is denoted by Y_T, where:

$$Y_T = 1,000 + 1,000 \max\left(g \ , \ \min\left(C, \frac{S_T - S_0}{S_0} \right) \right) \tag{1}$$

[7]Of course, in practice the sales agent may not provide the consumer with this choice but push the LC contract. Our comparison is still a valid exercise in order to better understand the market; and, of course, it is useful for policymakers and regulators to have this type of analysis available.

The initial investment of \$1,000 yields a rate of return between g and C at the end of the period. The payoff can be decomposed into a long position in a bond, a long position in a standard call option and a short position in another standard call option. More precisely,

$$Y_T = \mathbb{B} + \frac{1,000}{S_0}\left(\mathbb{C}_1 - \mathbb{C}_2\right) \tag{2}$$

where $\mathbb{B}$ is a bond that pays $1,000(1+g)$ at time T, $\mathbb{C}_1$ is a call option on S, with strike $S_0(1+g)$ and $\mathbb{C}_2$ is a call option on S, with strike $S_0(1+C)$. Both calls have the same maturity T. The globally-capped contract has a simple design because it can be replicated in terms of standard European options.

We next consider a quarterly-capped contract with a cap of c on the return over each quarter. Assume the minimum guaranteed rate is again g. Let $t_0 = 0$, $t_1 = \frac{1}{4}, t_2 = \frac{2}{4}, ..., t_{20} = T = 5$. The payoff (at maturity) X_T under the LC design can be expressed as:

$$X_T = 1,000 + 1,000 \max\left(g\ ,\ \sum_{i=1}^{20} \min\left(c, \frac{S_{t_i} - S_{t_{i-1}}}{S_{t_{i-1}}} \right)\right) \tag{3}$$

In the sequel, this contract is referred to as the *Quarterly Sum Cap*. It is a complex contract since the payoff is highly path-dependent and cannot be easily replicated with standard options.

3.2 Puzzle

There is empirical evidence that the complicated contract X is often subject to higher fees and yet the simpler contract Y can be easily replicated and should have lower fees. However, the puzzle is that the more complicated locally-capped contracts are very popular. This can be seen for instance from Figure 2. This figure summarizes the percentage in volume of locally-capped, globally-capped and contracts with no cap among the AMEX listed contracts that include a floor guarantee.

Figure 2: Exchange Listed Structured Products with Guarantee
We classified the structured products available on AMEX as of October 2006. We only include the products that provide a guarantee at maturity (about 50% of the total volume of exchange listed products). We split them into three categories: those with no cap, those with a local cap and those with a global cap. Most of the contracts with no cap are participating policies where the investors receive a percentage of the index return. The proportion in the pie chart corresponds to the percentage in volume of such contracts.

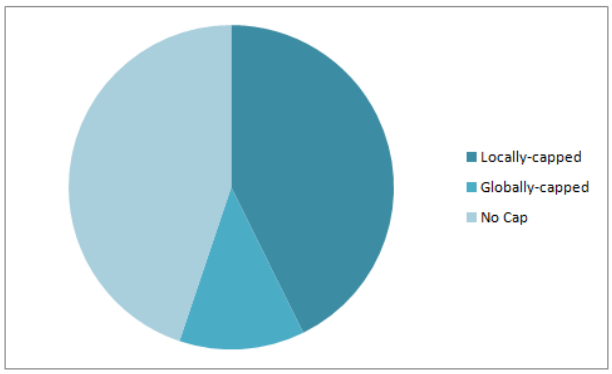

To make the comparison straightforward, we ignore the existence of higher fees and assume that both are fairly priced with a fixed commission of 8%. This means, that under the assumed market dynamics, the fair price of both contracts X and Y is equal to the initial investment of $x_0 = \$1,000 - \$80 = \$920$. We will show that the standard finance models based on the expected utility of final wealth fail to explain the demand and generally predict that consumers will prefer the simpler contract contrary to what is observed in practice.

3.3 Comparing Contracts with Different Cap Periods

We now describe the framework we use to compare two contracts which have different cap frequencies. We use concrete examples to facilitate the discussion. We use the Black and Scholes framework to model the equity portfolio return. The market is complete and arbitrage-free. We assume there are no transaction costs and no trading constraints. The underlying index is denoted by S and follows a geometric Brownian motion,

$$\frac{dS_t}{S_t} = (\mu - \delta)dt + \sigma dW_t$$

where W is a standard Brownian motion under the physical measure $\mathcal{P}$. In this model, we define the current *cost*, equivalently the *market value* or *no-arbitrage price* of the payoff X_T payable at time T as:

$$c(X) = E\left[\xi_T X_T\right] \tag{4}$$

where ξ_T is the state-price process.

We focus on a quarterly sum cap which has an initial value of \$920. Consider a five year maturity contract. In this case, the initial investment is \$1,000 and the guarantee is \$1,100 after five years. Under our assumptions the quarterly cap level works out to be 8.7% and this contract has an initial fair price of \$920. Theoretically, if the return of the index is greater than 8.7% for the twenty consecutive quarters, then the return of the contract is 174%. As one would expect, this event has an extremely low probability of occurring.

We wish to compare the quarterly-capped contract (that pays X_T) with the globally-capped contract (that pays Y_T) where both contracts are fairly priced and the fair price of each contract is 920\$. To do so, we find the global cap level for the five year contract that makes the fair price of the globally-capped contract equal to 920\$. The global cap level in this case is 30.5%. Thus, a contract with a global cap level of 30.5% has the same value as a contract with a quarterly cap level of 8.7%. Using the same approach, we can find the fair cap level for other contracts with different cap frequencies. These contracts could be yearly capped, semi-annually capped, monthly capped or even weekly capped. We can determine the caps of all these contracts such that they are all worth \$920 at inception. Note that all these contracts have the same global floor rate $g = 10\%$.

Figure 3: Equivalent cap levels of fairly priced contracts.
The horizontal axis represents the cap frequencies. The cap frequency ranges from weekly to five years (global cap). This frequency is quarterly for contract X and five years for contract Y. Parameter values are $r = 5\%$, $\delta = 2\%$, $\sigma = 15\%$. The quarterly cap is 8.7%. The corresponding *equivalent* global cap over 5 years is 30.5%.

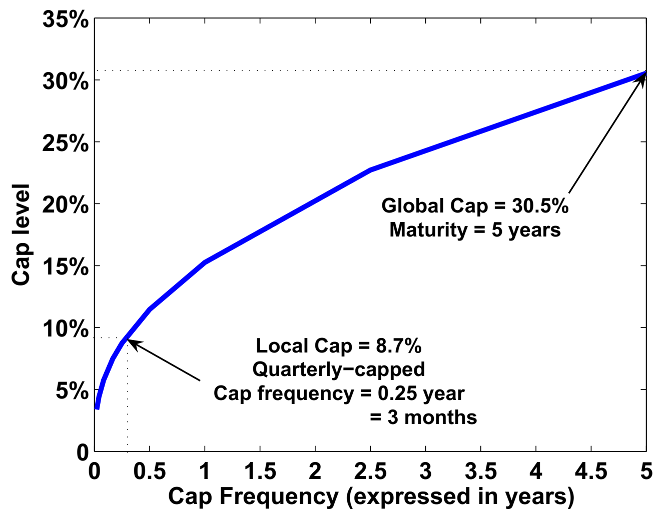

Figure 3 shows the cap levels corresponding to different cap periods. For each cap period, the cap level that gives the same fair value of the contract is computed. Table 2 summarizes the numerical values corresponding to Figure 3.

Table 2: Cap levels corresponding to different cap frequencies for contracts with the same market value equal to 920$ (Parameters are set to $r = 5\%$, $\delta = 2\%$, $\sigma = 15\%$).

Cap period	1 week	1 month	2 months	quarter	6 months	1 year	5 years
Cap Level	3.3%	5.8%	7.5%	8.7%	11.5%	15.2%	30.57%

From now on, when we discuss contracts with different cap frequencies, it is assumed that their market values are identical.

3.4 Path-dependent contracts are not cost-efficient

In this section, we use Dvbvig's (1988a, 1988b) approach to analyze the cost efficiency of the LC and GC contracts. Agents' preferences are assumed to depend only on the probability distribution of terminal wealth and they prefer more to less. Dybvig developed the approach using the Cox Ross Rubinstein binomial model. We use the Black and Scholes framework to model the equity portfolio return. In the Black and Scholes model, the state-price process can be written as $\xi_T = a\left(\frac{S_T}{S_0}\right)^{-b}$ where the parameters a and b are as follows

$$a = \exp\left(\frac{\theta}{\sigma}\left(\mu - \delta - \frac{\sigma^2}{2}\right)T - \left(r + \frac{\theta^2}{2}\right)T\right), \quad b = \frac{\theta}{\sigma}, \; \theta = \frac{\mu - \delta - r}{\sigma}.$$

We see that ξ_T is a function of the asset price S_T.

A strategy (or a payoff) is cost-efficient if it is not possible to construct another strategy that generates the same distribution at maturity but at a lower cost. Given F, a cumulative distribution function (cdf), we define its *distributional price* as:

$$P_D(F) = \min_{\{Y \mid Y \sim F\}} \{c(Y)\} \tag{5}$$

where $\{Y \mid Y \sim F\}$ is the set of random variables with cdf F. $P_D(F)$ is the cheapest way of generating the distribution F at maturity. The *efficiency cost* of a strategy with payoff X_T at T with cdf F is defined as:

$$\text{Efficiency cost of } X_T \;=\; c\left(X_T\right) - P_D\left(F\right) \tag{6}$$

Note that the cost c as well as the distributional price P_D depend on the state price process ξ. However for the ease of exposition, we omit ξ in the formula. A payoff is cost-efficient if and only if its efficiency cost is equal to zero.

Given a random variable X with cdf F. Its inverse F^{-1} is defined as $F^{-1}(y) = \min\{x \;/\; F(x) \geq y\}$. We will make use of the following result.

Proposition 3.1. *Let F be a cdf. Define*

$$X_T^* = F^{-1}\left(1 - F_\xi\left(\xi_T\right)\right) \tag{7}$$

then, the cdf of X_T^* is F. Moreover, for any random variable X_T with cdf F,

$$c\left(X_T^*\right) \leqslant c(X_T).$$

If in addition $c\left(X_T^*\right) < +\infty$, then X_T^* is the unique cost-efficient payoff that has the same distribution as X_T. Thus, it satisfies:

$$P_D(F) = c(X_T^*).$$

The proof is given in Appendix B.1. This result can be interpreted as follows. Without changing the distribution of the payoff at maturity, one can reduce the cost of the strategy by replacing the payoff X_T by X_T^* but keeping the same distribution under the physical measure and thus the same expected utility. For any utility function U, X_T and X_T^* are on a iso-utility curve (that is $E[U(X_T)] = E[U(X_T^*)]$ for all utility functions U).

As a consequence of this result path-dependent derivatives cannot be cost-efficient. Indeed, for each distribution of the final payoff, there exists a unique cost-efficient strategy that generates a given distribution. Thus, to be cost-efficient, the payoff of the derivative has to be of the form (7). It is a European derivative written on the stock S_T since the state-price process can be expressed as a function of S_T. We also note that the payoff X_T^* is a non-decreasing function of S_T. Another way to interpret the proposition 3.1 is to say that the necessary and sufficient condition in the Black and Scholes model for a payoff to be cost-efficient is for the payoff to be a non-decreasing function of S_T.

As a consequence of the preceding results, the locally-capped contract is not cost-efficient. Its efficiency cost is strictly positive. For our benchmark case, we assume the instantaneous return is $\mu = 0.09$, the volatility σ is initially set at 15%, the interest rate is constant and equal to $r = 5\%$ and that there is a continuous dividend yield $\delta = 2\%$. To to apply the above proposition, we need the distribution function of the locally-capped contract. We first compute this distribution numerically and then derive its efficiency cost.

Distribution of the locally-capped contract

Figure 4 displays the cumulative distribution function of the capped quarterly return of the index.

Figure 4: Single period (quarterly) return when there is an 8.7% cap.

Assume R denotes the quarterly return ($R = \frac{S_{t_1} - S_0}{S_0}$). The graph shows Pr(R $\leqslant$ x) where x is the horizontal axis. There is probability mass of 0.18 at 8.7% corresponding to the cap. Other parameters are: $r = 5\%$, $\delta = 2\%$, $\mu = 9\%$, $\sigma = 15\%$.

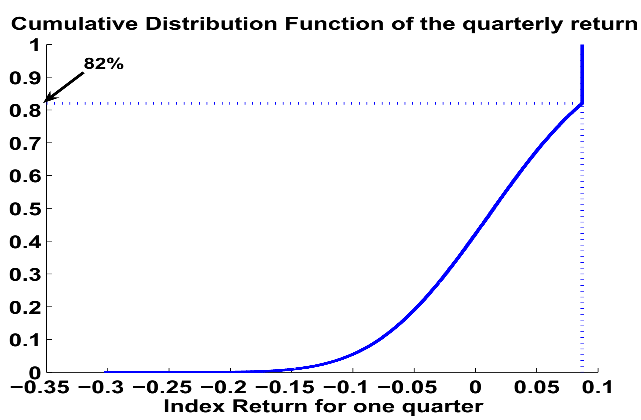

Based on our market assumptions, there is a probability mass of 18% at the cap level. As the number of quarterly returns in the sum increases, the left tail becomes more and more important and the maximum possible return increases by the value of the cap each period, with a diminishing mass of probability at the right end. The sum of 20 quarterly returns will be equal to 174% if and only if all consecutive quarterly returns exceeded 8.7%, so that the probability of this event is $0.18^{20} \approx 1.2 \times 10^{-15}$. In other words, this event is virtually impossible as can be seen from the cumulative distribution function of the locally-capped contract displayed in Figure 5. From Figure 5, it is clear that with probability almost equal to 1, the investor will get less than \$1,700. Under our assumptions, the Quarterly Sum Cap has a very high probability of yielding a return equal to ten per cent (with a probability mass there of 37%). It corresponds to the minimum guaranteed rate of ten per cent over the five year period.

If an investor is considering such a contract, increasing her awareness of the existence of the maximum possible return of 174% may influence her. We provide precise examples in section 5 to illustrate current practice on the AMEX and how sales agents may draw attention to this return and in so doing may influence the investor.

Figure 5: Cdf of the LC contract when there is an 8.7% quarterly cap. Now we study the cdf of the payoff of the LC contract. The graph shows $\Pr(X_T \leqslant x)$ where x is the horizontal axis. There is probability mass of 0.37 at 1100 corresponding to the guarantee at maturity. Other parameters are: $r = 5\%$, $\delta = 2\%$, $\mu = 9\%$, $\sigma = 15\%$.

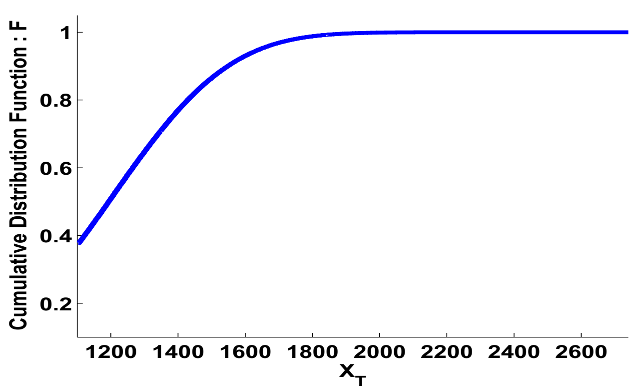

Efficiency cost

Given these parameters, the efficiency cost of the locally-capped contract is \$1.1 when the cost (market value) of the contract is $c(X_T) = \$920$. It is dominated by a payoff X_T^* whose expression is given by (7) where F is the distribution of the payoff X_T (F is represented in Figure 5) and F_ξ is the cdf of the state-price process ξ_T. The payoff X_T^* is distributed as the locally-capped contract but has a lower cost, its cost is $c(X_T^*) = \$918.9$. A numerical demonstration of formula (7) is displayed in Figure 6. It is interesting to know that the LC contract is dominated by the payoff represented in Figure 6 in the expected utility framework for all increasing utility functions. This cost-efficient payoff is slightly concave and capped at \$2,740 since it has the same distribution as the LC contract. From Figure 6, we can see that high payoffs (say, over \$2,000) occur when the stock price after 5 years is approximately over \$600 (which has a probability of occurring virtually equal to 0). Note also that the probability to obtain the guarantee with the payoff X_T^* is also 37%. It is the probability that the stock price is below \$150 after 5 years (when $\sigma = 15\%, \mu = 9\%, \delta = 2\%, r = 5\%, S_0 = 100$).

Figure 6: Cost-efficient point-to-point design with same distribution as the LC contract
With $\sigma = 15\%, \mu = 9\%, \delta = 2\%, r = 5\%, S_0 = 100$, $T = 5$ years, quarterly cap $c = 8.7\%$ the graph represents the payoff of the cost-efficient derivative (as a function of the stock price S_T) that has the same distribution as the Quarterly Sum Cap.

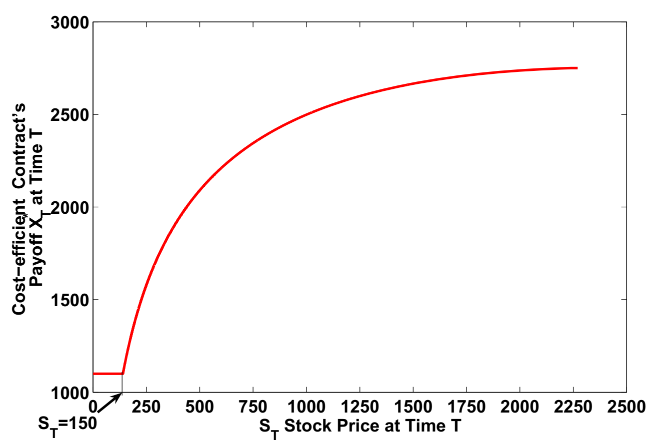

Note that a globally-capped contract is cost-efficient since its payoff is a non-decreasing function of the stock price at T. Its efficiency cost is equal to zero. This approach allows us to identify payoffs that are not optimal, such as the locally-capped contracts. If the decision criteria of the investor is based on the efficiency cost, then the GC design should be preferred. Note that it is not possible to directly compare the GC design and the LC design in terms of cost effectiveness since these two contracts do not have the same distribution at maturity T. Thus they will not have the same expected utility for all utilities. The following approach allows us to make comparison directly between the two designs but we have to make specific assumptions about investor preferences.

3.5 Choosing between Different Contracts: A Puzzle

We now consider the decision of an investor who is selecting between a locally-capped contract and a globally-capped contract. We use standard finance tools to examine this choice. First, we assume the investor has mean variance preferences and then investigate a risk averse investor with exponential utility.

3.5.1 Mean-Variance Criteria:

In this section we assume the investor has mean variance preferences. We will use the excess expected return over the guaranteed return to compute a *modified Sharpe ratio*[8] to make the comparison. Let Z_0 be the initial investment and let the guarantee be $(1+g)Z_0$ at the maturity T. We define the modified Sharpe ratio as follows

$$R_Z = \frac{\mathrm{E}[Z_T] - Z_0(1+g)}{\mathrm{std}(Z_T)},$$

[8] We will use the term Sharpe ratio rather than modified Sharpe ratio when there is no chance of confusion.

where the expectation is taken over the investor's own subjective probability measure which we assume to be the physical measure $\mathcal{P}$. We subtract the value of the guarantee to ensure the Sharpe ratio is positive since the expected return of the contracts X and Y can be lower than the risk-free rate (because of the presence of 8% commission at inception of the contract). In particular, the expected return of the Quarterly Sum Cap is often below the risk-free rate when the volatility is high and the usual version of the Sharpe ratio becomes negative. We compute this ratio for the quarterly-capped contract R_X and for the globally-capped contract R_Y. The expectation and the standard error are calculated using the Monte Carlo method.

Figure 7: Comparison of Modified Sharpe Ratios.

The Quarterly Sum Cap has a quarterly cap of 8.7%, a global floor $g = 10\%$ and a maturity $T = 5$ years. We denote the modified Sharpe ratios by R_X (quarterly-capped contract) and R_Y (globally-capped contract) for different volatility levels. For each volatility, the cap of the globally-capped contract is such that the contract has the same no-arbitrage price as the benchmark 8.7% quarterly-capped contract. Other parameters $r = 5\%$, $\delta = 2\%$, $\mu = 0.09$.

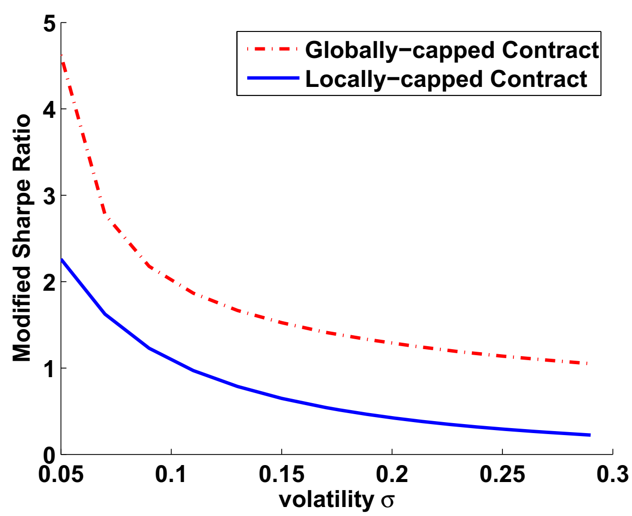

Figure 7 illustrates the comparison. It shows that the Sharpe ratio of the globally-capped contract exceeds that of the quarterly-capped contract for all volatility levels. This result indicates that a mean-variance investor should always choose the simple contract with the global cap rather than the complex contract with the quarterly cap.

The lower the volatility of the underlying index, the higher are the Sharpe ratios. Note that a lower volatility increases the value of the Quarterly Sum Cap. However recall that the two contracts are fairly priced and that both contracts have exactly the same no-arbitrage price, so that the contract price does not influence the comparison.

3.5.2 Other Risk Preferences

We now conduct the same comparison using more general risk preferences to see if the results for the mean variance preferences still hold. We illustrate our results with exponential utility preferences but similar conclusions are available for other concave utility functions. We will see that for some parameter values the investor's prefers a contract with a lower cap frequency to one with a higher cap frequency. For other parameter values the opposite is true. Our analysis shows that the key factors include the investor's degree of risk aversion and the assumed volatility level. We can roughly summarize our findings as follows by taking the quarterly sum cap and the global cap as the two reference contacts. Other things equal the Quarterly Sum Cap becomes more attractive as the investor becomes less risk averse. Other things equal, the lower the assumed volatility the more attractive the Quarterly Sum Cap becomes.

We use the exponential utility assumption:

$$U_a(X) = \frac{1 - e^{-ax}}{a}.$$

where a is the *coefficient of absolute risk aversion.* Figure 8 consists of two panels to illustrate the impact risk aversion. Panel A corresponds to an investor with a very low absolute risk aversion: $a = 0.005$ and Panel B corresponds to a very risk averse investor with an absolute risk aversion: $a = 8$. In both panels, the expected exponential utility is calculated for two levels of volatility $\sigma = 10\%$ and $\sigma = 20\%$. The horizontal axis corresponds to the cap frequency. The cap frequency varies between one week and 5 years. We again make the assumption that all contracts have the same no-arbitrage price equal to \$920.

Figure 8: Impact of Risk Aversion & Index Volatility
These results are based on the following parameters: $r = 5\%$, $\delta = 2\%$, $\mu = 0.09$. $g = 10\%$, $T = 5$ years. In Panel A, the coefficient of absolute risk aversion is $a = 0.005$ corresponding to an investor with a very low risk aversion and in Panel B, $a = 8$ corresponds to a very risk-averse investor.

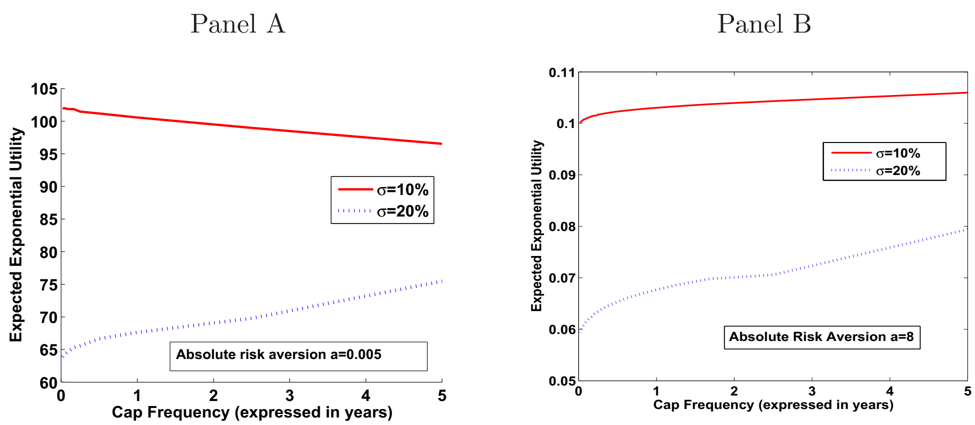

Panel A in Figure 8 shows that investors with low risk aversion will prefer the Quarterly Sum Cap to the globally-capped contract when the volatility is low (10%). However when the volatility is 20% this preference is reversed and the globally-capped contract is preferred. Panel B shows that risk averse investors will prefer the global cap to the local cap under both volatility assumptions.

4 Discussion and Possible Explanations

We showed in the previous section that the locally-capped contract will never be optimal in the expected utility framework and will rarely be preferred to a consistently priced globally-capped contract. There is a large volume of contracts with local caps both in the AMEX listed securities and in the EIA market (see Appendix A for an overview of locally-capped globally-floored products).

The relative popularity of locally-capped contracts cannot be readily explained within the expected utility framework. One possible explanation is to assume that investors have more complex decision criteria. Since locally-capped contracts are path-dependent, one could envisage path-dependent preferences. However, these contracts are European contracts and most of them have no cash-flows prior to maturity. Thus, it is hard to make the case for path-dependent preferences. Another approach is to revisit how individuals subjectively assess risks. Subjective probabilities can be very different from true probabilities and this could play an important role in explaining the demand for locally-capped products.

We first recall the main ideas of Kahneman and Tversky (1979) and explain how they can apply to our framework. We then analyze more specifically why the payoffs of the quarterly-capped contracts are difficult for retail investors to understand. Due to their complexity and the influence of the sales agents, retail investors might overweight the probability of getting high returns. We argue that the demand for locally-capped products can be explained by this overweighting. This misrepresentation of the returns is encouraged by the hypothetical examples presented in sales prospectus (such as the Example 1 reported in appendix C). We discuss this point more fully in section 5.

We also consider another possible explanation in subsection 4.4 proposed by Carlin (2008). Due to the high complexity of the product, investors may choose to stay uninformed and select a product randomly among the menu of available contracts. They can easily pick the complex product that is not optimal for them. If the sales agent receives a higher commission on selling the complex contract, this will exacerbate the situation.

4.1 Beyond the Standard Concave Utility Framework

Kahneman and Tversky (1979) argued that agents' attitudes towards risks are different for gains and losses. Their theory has several aspects: a kink point at the origin, a convex utility function over losses, a concave utility function over gains and an overweighting of tail probabilities. The kink point at the origin was introduced to explain the fact that agents are more sensitive to losses than to gains. Indeed, Kahneman and Tversky (1979) observed high risk aversion against asymmetric bets such as for instance a 50% chance of a gain of 1,100 and a 50% chance of a loss of 1,000. Then, as noted also by Barberis and Huang (2007) most people prefer a certain gain of 500 over a gain of 1,000 with probability 1/2. In contrast, they prefer a loss of 1,000 with probability 1/2 to a certain loss of 500. This explains the fact that individuals are risk averse over moderate probability gains and risk-seeking over moderate-probability losses.

This behavior may not be representative when larger amounts of money are involved as shown by the following example. An individual may well prefer to bet on an investment that can give them \$5,000 with a probability 0.001 rather than a certain gain of \$5. This can be explained by assuming that investors overweight tail probabilities, in particular they may overweight the probability of getting \$5,000. Under prospect theory, investors do not use true probabilities when evaluating investment decisions but rather overweight probabilities. When making their decisions investors assume $\pi(\$5,000) > 0.001$. Thus, they prefer the bet rather than the certain gain. Barberis and Huang (2007) mention that "the transformed probabilities ... should not be thought of as beliefs, but as decision weights that help capture evidence on individual risk attitudes". They also note that in Kahneman and Tversky's

(1979) framework, investors know the true probabilities and "the overweighting of 0.001 introduced by prospect theory is simply a modeling device".

Our problem is slightly different but a partial explanation can be derived from these ideas. In our case investors do not have a perfect assessment of the true probabilities. We make use of overweighting techniques, not only as a modeling device but interpret them as biased beliefs. We document in section 5 why investors would misrepresent the possible outcomes of investment returns and overweight probabilities of getting high returns because both of the complexity and of the available information provided by sales agents.

4.2 Distribution of the Payoff of a Quarterly Capped Contract

In this section, we come back to the analysis of the return distribution under a Quarterly Sum Cap. We highlight specific features that illustrate the complexity of the product. We also suggest why investors may form a biased view of the payoff distribution and overestimate the returns on locally-capped contracts.

The distribution of the payoff of a Quarterly Sum Cap is quite complicated and it is extremely difficult for the typical retail investor to have a realistic picture of this distribution. The reason stems from how the cap affects the final distribution of returns. Assuming that that there is *no* cap, investors might have a reasonable idea of the distribution of the sum of independent and identically distributed quarterly returns. But because of the presence of a cap, the return $\frac{S_{t_{i+1}}-S_{t_i}}{S_{t_i}}$ each quarter has a truncated distribution function as shown in section 3 in Figure 4. Based on our benchmark assumptions, the expected quarterly return of the underlying index is 1.8%. After truncation at the cap level, the expected quarterly return declines to only 0.9%. If the investor does not really understand the impact of the local cap on the quarterly return, she might think the average of the sum of 20 returns is about $0.018 \times 20 = 36\%$ instead of only 18% ($0.009 \times 20 = 18\%$). Note that 18% return over

5 years corresponds to a 3.3% continuously compounded annual rate which is lower than the risk-free rate of 5%.

In section 3, we already displayed the cdf of the capped quarterly return of the index (see Figure 4). Based on our market assumptions, there is a probability mass of 18% at the cap level. As the number of quarterly returns in the sum increases, the probability of reaching the cap level each and every quarter declines very quickly. For instance after three quarters, the maximum value for the sum of the three returns is equal to 26.1%, which means each of the three returns is equal to the cap level of 8.7%. If the quarterly returns are independent the probability of getting the cap level each quarter is simply equal to $0.18^3 \approx 0.006$. Note that the sum of 20 quarterly returns will be equal to 174% if and only if all consecutive quarterly returns exceeded 8.7%, so that the probability of this event is $0.18^{20} \approx 1.2 \times 10^{-15}$.

If an investor is considering such a contract, increasing her awareness of the existence of the maximum possible return of 174% may influence her. It is very likely that the investor overweights the probability of this maximum return over and beyond the *natural* overweighting of rare events that has been documented in the literature. The sales agent may draw attention to this return and in so doing may influence the investor. Numbers and examples that appear in a prospectus or sales literature can play a role in forming expectations. We provide precise examples in section 5 to illustrate current practice on the AMEX.

In both contracts that we examine, there is a minimum guaranteed rate of 10%, which corresponds to a global floor over the five year period. Figure 9 shows the probability distribution functions of the payoffs of a Quarterly Sum Cap, X_T, and the globally-capped contract, Y_T. Under our assumptions, the Quarterly Sum Cap has a very high probability of yielding a return equal to 10% (with a probability mass there of 37%). The probability mass at 174% in the right tail is 1.2×10^{-15}. The globally-capped contract Y has a global cap rate of 30.5% and the same floor rate of 10%. Hence this distribution has *two probability* *masses*: one at 10% with probability 27.7% and one at 30.5% with probability 53%. Thus, the investor has a probability of only 19.3% of obtaining a return between these two numbers. The returns are almost uniformly distributed in the continuous part of the distribution (as shown by the graph on the right in Figure 9).

Figure 9: Comparison of Probability Distribution Functions.
The left panel is the density of the maturity payoff under the Quarterly Sum Cap (X). The right panel corresponds to the density of the maturity payoff under the globally-capped contract (Y). Input parameters are: $r = 5\%$, $\delta = 2\%$, $\mu = 0.09$, $\sigma = 15\%$.

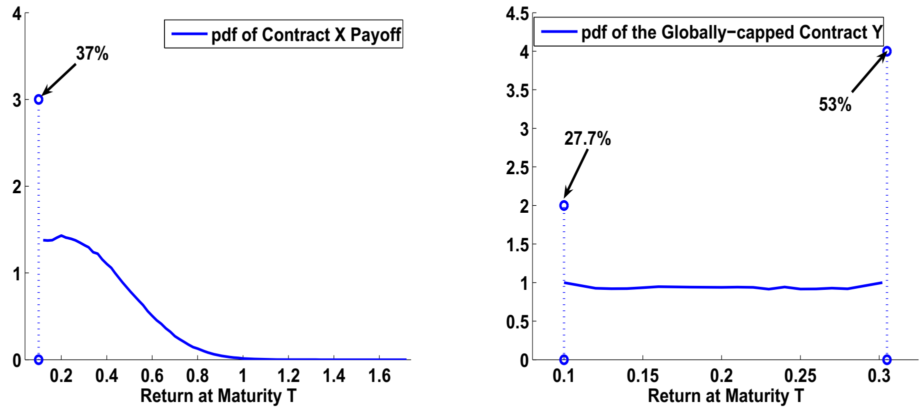

4.3 Biased beliefs

As mentioned before, sales agents can draw attention to the maximum attainable return[9]. Thus, it is extremely difficult for retail investors to have an accurate notion of the distribution of these payoffs under different contracts. In the Quarterly Sum Cap example, the maximum return is 174% over one year. By stating this fact, investors may be encouraged to believe that this could really happen. The consumer will have biased beliefs on the distribution of possible payoffs and will overweight the probabilities of getting high returns.

There are many precedents in the literature that support this approach. This idea is used

[9]See section 5 for empirical support of this claim.

for instance in cumulative prospect theory in Tversky and Kahneman (1992). Brunnermeier, Gollier and Parker (2007) consider optimistic agents and show that "an investor chooses to be optimistic about the states associated with the most skewed Arrow-Debreu securities that are either the least expensive state (when states are equally likely) or the least likely state (when state prices are actuarially fair)". Thus, they consider subjective probabilities. The concept dates back Ali (1977) or Machina (1982). Ali (1977) shows that if a horse has a high objective probability of winning a race then punters will tend to understate the probability of winning. Conversely, if a horse has a low probability of winning punters will tend to overstate its probability of winning.

In our case, the probabilities of getting high returns under the Quarterly Sum Cap are very low, but investors may believe they are more likely than they really are, and overweight the right tail of the distribution. We do not have a closed-form expression for the density of the sum of the twenty quarterly-capped returns and we have a mixture of continuous distributions and discrete distributions. Both of these features make it difficult to apply standard overweighting techniques (see Appendix B.3 for more details). We adopt a very simple approach that is easy to implement and has a straightforward economic intuition.

We propose to overweight the right tail of the distributions by combining two techniques. First, we increase the drift of the underlying index, which gives a simple practical way of overweighting the entire right tail of the distribution. It has a simple intuitive meaning in that it is equivalent to assuming the expected return on the index increases by a certain percentage. Second, we also add a lump of probability at the extreme right end of the distribution. The probability of getting a 174% return for the Quarterly Sum Cap is 1.2×10^{-15} in reality, but in the mind of investors, it is a possible event and they thus may assign a higher subjective probability to this event. We assume their subjective probability of this event is increased to 1%. We thus increase the probability of getting 174% over five years by 1% (and lower the probability of getting the guarantee by 1%).

Changing the drift of the underlying is equivalent to a change of measure using Girsanov theorem (see Appendix B.3 for details). This is a simple adjustment to the original probability measure. The cap level will be reached more frequently than under the original measure. The net result is thus to overweight the right tail of the distribution and to underweight probabilities in the left tail, as for instance the probability of getting the guarantee at maturity[10].

4.3.1 Illustration of Overweighting the Right Tail

We now illustrate the impact of overweighting the right tail of the distributions of the payoffs under the Quarterly Sum Cap and the globally-capped contract (as we did in Figure 9). The parameters are unchanged except that the investor has now a biased representation of the probability distribution of the Quarterly Sum Cap. We leave the probability distribution of the globally-capped contract unchanged since it is truncated at both ends and the distribution does not have any tails.

Under our assumed real world probability measure, the Quarterly Sum Cap has a very high probability of yielding 10% (with a probability mass 37%) and a very small probability mass of 1.2×10^{-15} at the right end of the tail corresponding to the maximum possible return of 174%. If the investor is more optimistic concerning the probabilities of high returns, she overweights the right tail by overestimating the expected return on the index. If we assume the additional return is 5% per annum, then the probability of obtaining 10% is only equal to 14% and of obtaining 174% is 1.3×10^{-13}. This event can still be considered as impossible event. Figure 10 displays the subjective distribution of an optimistic investor

[10]Under the risk neutral probability, the drift is r instead of μ, and the no-arbitrage prices of a LC contract and a GC contract are identical. Using the risk neutral measure underweight the right tail. The expectation under the actual probability is higher then for the LC contract (explaining the preference for quarterly sum caps by risk neutral agents).

who overestimates the yearly index return by 5%. This modification shifts weight from the mass at the guarantee to the right.

Figure 10: Overweighted Probability Distribution Functions.
We display the density of the payoff under the Quarterly Sum Cap (X) with an additional expected annual Index return of 5%. The quarterly cap is $c = 8.7\%$. Other parameters : $r = 5\%$, $\mu = 9\%$, $\delta = 2\%$, $\sigma = 15\%$.

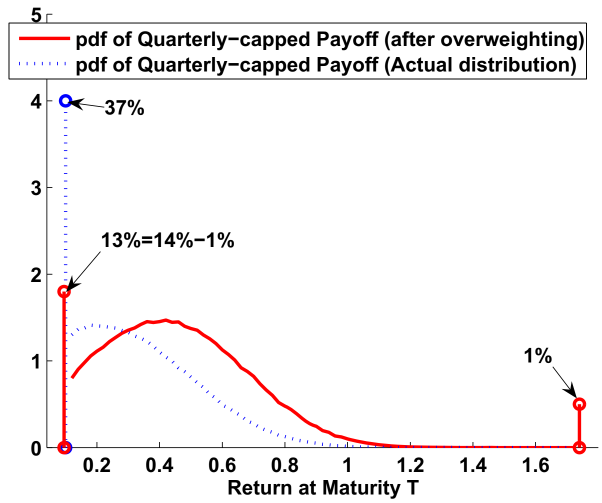

4.3.2 Investor Choice under the New Probability Measure

We now analyze the impact of our proposed overweighting on the relative preference for the Quarterly Sum Cap and the Globally-Capped Contracts.

We apply the same decision criteria as before to see what difference the overweighting makes to the investor's contract choice. First we compute the modified Sharpe ratio using the new measure for the Quarterly Sum Cap and the original measure for the other contract. The formula is

$$\tilde{R}_X = \frac{\mathrm{E}_Q[Z_T] - Z_0(1+g)}{\mathrm{std}_Q(Z_T)}$$

We then compare this Sharpe ratio $\tilde{R}_X$ with R_Y that we computed earlier. Figure 11 illustrates the result. In Panel A, investors overweight the right tail of the distribution by increasing the expected return on the index by either 3% or 5%. In Panel B, the investor also believes that there is a additional probability mass of 1% at the maximum possible return 174% for the Quarterly Sum Cap distribution. We see that the Quarterly Sum Cap is preferred to the global contract when the volatility is low. This effect is accentuated when there is an additional lump of probability in the right hand panel at the maximum possible return.

Figure 11: Sharpe Ratios when returns are overweighted.

The figure shows the Sharpe ratio $\tilde{R}_X$ (quarterly-capped contract), R_Y (globally-capped contract) for different volatility levels. The original contract has a 8.7% quarterly cap, $g = 10\%$, $T = 5$ years. For each volatility the cap of the globally-capped contract is such that the contract has the same no-arbitrage price as the 8.7% quarterly-capped contract. In Panel A, investors overweight the tail of the distributions by increasing the annual expected return by either 3% or 5%. Panel B, in addition has a lump of 1% at the right end of the tail when the return is maximum, that is 174%. Other parameters $r = 5\%$, $\delta = 2\%$, $\mu = 0.09$.

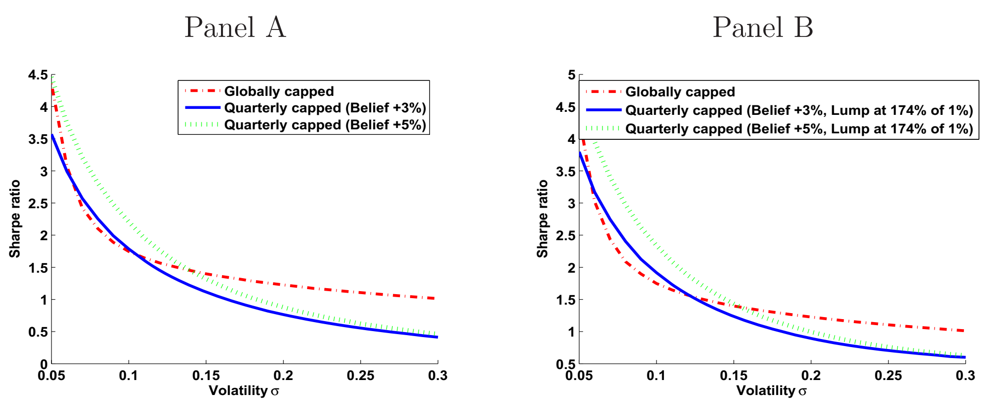

Recall that when we used the original measure and did not overweight the probabilities a mean variance investor would always prefer the globally-capped contract to the Quarterly Sum Cap. When we overweight the probabilities this preference is reversed for low volatility assumptions. In order for investors to prefer the quarterly sum cap to the globally-capped, say when the volatility $\sigma = 15\%$, they have to increase the annual expected return by 5%

p.a. or to increase the annual expected return by 3% p.a. but to believe there is a probability of 1% of getting the maximum possible return.

The second criteria is based on the exponential utility. The absolute risk aversion is $a = 0.005$ in the left panel of Figure 12 and $a = 8$ on the right panel as we did in Figure 8. Figure 12 compares the expected utility of the two contracts for different levels of overweighting when the volatility is 20%. Recall that for both values of the absolute risk aversion, it was not optimal to pick the quarterly sum cap when the volatility was 20%. In Panel A, when the absolute risk aversion is 0.005, it is optimal to pick the globally-capped contract when investors overestimate the expected annual return on the index by about 8% per annum. In Panel B, overweighting the probability distribution by 12% per annum is still not enough to make the Quarterly Sum Cap more attractive than the globally-capped contract. When the volatility is high and investors are risk-averse, the globally-capped contract is preferred even if investors significantly overestimate high returns.

Figure 12: Sensitivity of Investor choice to extent of overweighting

Expected Exponential utilities of two contracts (Locally-capped and Globally-Capped) under different overweighting assumptions. $r = 5\%$, $\delta = 2\%$, $\mu = 0.09$, $\sigma = 18\%$. The overweighting is done through the belief of an additional return over μ. A lump of 1% is added at 174%. Here the absolute risk aversion is a. Panel A: $a = 0.005$, Panel B: $a = 8$.

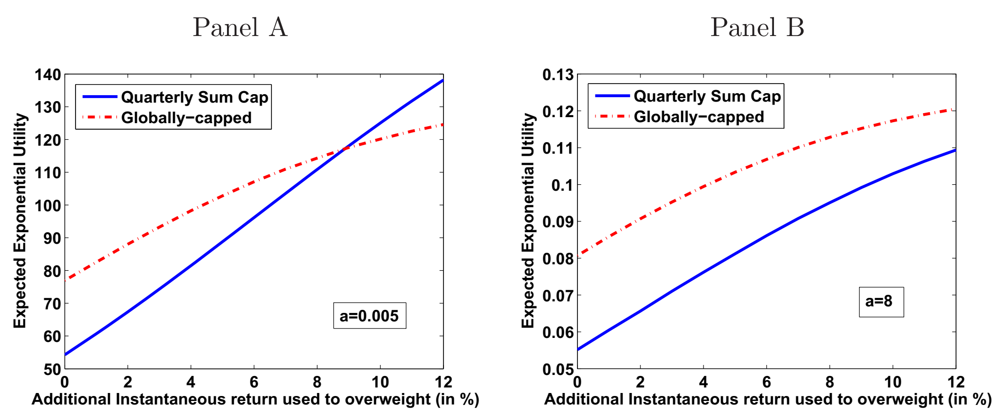

Overweighting the right tail can induce investors to prefer the Quarterly Sum Cap to the globally-capped contract and hence this provides a possible explanation for the popularity of locally-capped contracts. We see from Figure 11, that effect is amplified when the volatility is low and thus when retail investors underestimate the volatility of the market.

An interesting feature of globally-floored contracts is that they are not subject to losses because of the presence of the minimum guaranteed rate. Even if the minimum guaranteed rate is zero, the initial investment is fully guaranteed. It is possible that when there are no losses investors might become risk seeking rather than risk averse. In this case, a convex utility function on a suitable range might also explain why investors prefer the locally-capped contracts.

4.4 Complexity in Structured Products

We now discuss the issue of product complexity in the case of retail structured products. As noted earlier, Carlin (2008) has developed a model where the sellers of retail financial products deliberately design them to be complicated in order to confuse consumers and increase profits. We will show that certain features of the structured products market are consistent with Carlin's model.

We begin by noting that, to the best of our knowledge, there is no clear definition of product complexity in the case of financial products. Nevertheless, there are many cases where most observers would agree that one contract is more complex than another. For example, a risk-free zero coupon bond would be universally viewed as being less complex than a locally-capped, globally-floored Equity Index Annuity just as a toothbrush is less complex than an iPod. In other cases, the rankings are less obvious. We suggest that the number of features in the contract, the difficulty of figuring out the benefits and the amount of technical jargon used to describe it would all appear to contribute to product complexity.

On this basis, a locally-capped contract is more complex than a globally-capped one.

It is hard to argue against the proposition that the typical structured financial product is quite complex. For example, we have shown that it takes considerable analysis to work out the payoff distribution of locally-capped globally-floored contracts. Of course, there could be welfare benefits for consumers, when producers experiment with new innovative designs that better complete the market. It is beyond the scope of this paper to provide an analysis of this trade-off. However, it is reasonable to suggest that structured financial products are unnecessarily complex. If this is accepted, then the market is consistent with Carlin's model.

Carlin also predicts that producers will increase the complexity of their financial products in order to overprice them. Again, we can find features of the structured products market that are consistent with this prediction. EIAs and exchange listed notes often provide similar investment benefits. The sales commissions on index listed notes are generally around 5% whereas the average EIA commission in 2006 was about 8% (Koco (2007)). Furthermore, the insurance industry has actively lobbied to avoid SEC regulation which would result in more disclosure which in turn would presumably lead to a reduction in complexity and commissions. The association between overpricing and product complexity could also help explain the popularity of the locally-capped globally-floored contracts we examined in this paper. In summary, the evidence from the structured products market, preliminary as it may be, confirms the connection between product complexity and overpricing.

5 How the investors probabilities can be overweighted

This section provides empirical evidence that as to how an investor's probabilities of high returns could be overweighted in the case of locally-capped contracts. We document examples from the various Pricing Supplements associated with the AMEX contracts.

Each of the three J.P. Morgan contracts we discussed earlier contains five hypothetical examples of future S&P performance and the corresponding payoffs to the investor. The first two examples of the product *JPL.G* are reported in Appendix C. The first example assumes that the quarterly return on the Index is 6% each quarter, while the other examples illustrate more realist forecasts. If the return on the Index is 6% each quarter, the investor will receive $2,200 per $1,000 of principal at maturity. Of course this event has an extremely low probability. The inclusion of this very optimistic scenario even as an hypothetical illustration may cause some investors to regard this event as being more probable. The Pricing Supplements for the other two tranches also contain five hypothetical examples. In each case, the first example also shows the index hitting the cap level each quarter for each of twenty successive quarters. These examples indicate how the natural overweighting of the probability of high returns can be increased. We have already shown that overweighting the probability of high returns will increase the attractiveness of the locally-capped contract versus a globally-capped contract. Thus, this provides a possible explanation for the popularity of the locally-capped contracts.

We found that in the case of all 39 locally-capped globally-floored products at least one projection assumed the cap level was attained every period. Our analysis of these 39 prospectuses reveals that the above description is common practice and that the JP Morgan illustrations are far from extreme. Indeed, our analysis reveals that all issuers provide in their prospectus 4 to 7 hypothetical examples. One or two of the first three examples assume that the investor receives the maximum possible return. In most of the cases, the first three examples are unrealistic hypothetical examples. These examples are based on very unrealistic projections and may bias investors perceptions. This is consistent with the way investors assess low probability events (Camerer and Kunreuther (1989)). In particular they note that adding details to the description of events make them appear to be more likely because such details add plausibility.

The most extreme *hypothetical examples* were observed in the case of compounded monthly-capped products (which are the products that provide highest possible returns). We now illustrate the most extreme set of unrealistic assumptions we found. This product is based on the Nasdaq under the name *NAS: Nasdaq-100 Index TIERS*. The initial investment is \$10 and the maturity payoff is a compounded monthly return capped at 5.5% per month. In the prospectus, there is a description of 7 hypothetical examples. In the first three examples, it is assumed that the Nasdaq-100 index increases by 3%, 5.5% and 7% during each and every one of the 66 months of the contract. Under these assumptions the final payoffs are respectively $1.03^{66} = \$60.35$, $1.055^{66} = \$332.5$, $1.055^{66} = \$332.5$. Examples 4, 5 and 6 display cases where investors receive the minimum guarantee based on a low interest, that is \$10.7. In the last example, the payoff is \$13.35.

Out of the 278 monthly returns on the Nasdaq between October 1985 and November 2008, there were 60 months when the monthly return exceeds 5.5%. On this basis the empirical probability of a monthly return exceeding 5.5% is thus $60/278 = .22$, say 0.2. Assuming an i.i.d. distribution of the monthly returns, the probability of experiencing the maximum possible return is $0.2^{66} = 7 \times 10^{-47}$ which is an impossible event. Thus the probability of examples 2 and 3 occurring is less than 10^{-46}. The probability of getting the guarantee is on the other hand very high as we noted in other examples. Historically, the probability of observing the returns in Examples 1, 2, 3 and 7 is zero. Indeed the maximum value for the compounded return of 66 consecutive monthly capped returns is 2.7 observed in May 1996. Examples 4 and 5 have an historical probability of about 50% of taking place.

In addition, these securities are subject to default risk. The prospectus states *"Principal-Protected" means that your principal investment in the certificates is not at risk due to a decline in the Index. "Principal-Protected" does not mean that you will receive a return of* *your principal investment in all cases. Under certain circumstances, losses realized on the assets of the trust will be borne by the holders of the certificates. In particular, upon the occurrence of a default by the swap counterparty and the swap insurer or upon the occurrence of a term assets credit event, you may receive less than the principal amount of your investment."* Until very recently, retail investors may not be aware that banks may go bankrupt. Since these securities are not protected by the FDIC investors could lose some or all of their principal.

6 Conclusions

This paper described some contract designs that are popular in the retail structured products market. We showed that some of the consumer choices made in this market are very puzzling from the perspective of standard finance theory. We analyzed locally-capped globally-floored contracts and compared them to globally-capped globally-floored contracts. The latter should be preferred under standard models of investment behavior. The locally-capped products have path-dependent payoffs and are inefficient contracts if the investor only cares about the expected utility of end of period wealth. Our paper demonstrates that this preference relation could be reversed if investors overweight the probability of high returns in the locally-capped contract. We showed that overweighting the tails of the distribution arises in a very natural way in locally-capped contracts and we provided evidence that this tendency is encouraged by the hypothetical projections in the prospectus supplements. It is hoped that this paper contributes to our knowledge of investor choice in the retail structured products market.

Since August 2008 there has been a massive increase in the volatility of financial markets and a worldwide credit crisis. These conditions will have some important impacts on the market for structured products. The implied volatility of the $S\&P$ has reached levels of over 60%. We note that from Figure 1 that at these assumed volatility levels the option component of the locally capped contract is worth almost zero. On the other hand, in times of poor stock market performance maturity guarantees become important. However if the entity providing the guarantee is subject to credit risk, the value of the guarantee may be illusory. Overall we would expect the demand for these products to decline in the current environment.

A Locally-Capped Globally-Floored on AMEX in Oct. 2006

There were 39 locally-capped globally-floored index linked notes corresponding to a total volume of \$2.6 billion. Here is the list of issuers of these products, we indicate the number of issued products in parenthesis.

- The Bear Stearns Companies (3),
- Structured Products Corp., Ambac (5),
- Citigroup Global Markets Holdings (5),
- Bank of America (12),
- Morgan Stanley (6),
- International Finance Corp. (1),
- J.P. Morgan Chase & Company (3),
- Lehman Brothers Holdings (1),
- UBS AG (3).

These products are mostly linked to US indexes in particular 14 of them are linked to the S&P 500 index, 5 of them to the Nasdaq 100 index and 9 to the Dow Jones Industrial Average. Table 3 summarizes our data:

Table 3: Overview of Locally-Capped Globally-Floored Products

These products are locally-capped (monthly or quarterly) and the adjusted return can be computed as the sum of the capped returns or the compounded product of the capped returns.

Adjusted Return	Number	Average Cap	Average Term	Average Guarantee
Monthly Sum	2	4.25%	5.25 years	100% of Principal
Compound Monthly	9	4.7%	5.3 years	100% of Principal
Quarterly Sum	8	6.5%	5.25 years	111% of Principal
Compound Quarterly	20	8.6%	5.3 years	103% of Principal

Note that highest returns are provided by compounded monthly capped returns. For example with a 4.7% monthly cap, the maximum yearly return of a compounded monthly capped return is $1.047^{12} = 1.73$ compared to $1 + 1.047 \times 12 = 1.564$. For a 7% quarterly cap, the highest yearly return for the compounded quarterly capped return is $1.07^4 = 1.31$ whereas the sum of quarterly capped return cannot exceed $1 + 0.07 \times 4 = 1.28$.

We noticed that the largest volume corresponds to 3 products guaranteed by Ambac for a total of \$785 million. Their names on AMEX are respectively *DJQ: TIERS Principal-Protected Minimum Return Trust Certificates (Interest based upon the Dow Jones Industrial Average) When Issued, NAS: Nasdaq-100 Index TIERS, Series 2003-13 When Issued, SPO: TIERS Principal-Protected Minimum Return Asset Backed Certificates Trust Series S&P 2003-23 When Issued*. These three index-linked notes have compounded monthly caps and their average maturity is just over 5.5 years.

B Technical Results

B.1 Proof of the proposition 3.1

Problem (5) defined by the distributional price can be reformulated as follows:

$$\min_Z E_P\left[\xi_T Z\right] \quad \text{subject to} \quad \forall x \in \mathbb{R}, \quad P(Z \leqslant x) = F(x) \tag{8}$$

The objective is to minimize the cost. The constraint is that Z has the same distribution as F (since the cdf of Z is F). This problem has been solved by Jin and Zhou (2008) [see Theorem B.1 (ii)]. The proof given by Jin and Zhou shows that for any feasible solution Z, that is for any random variable that is distributed with the cdf F one has:

$$E\left(\xi_T Z^*\right) \leqslant E\left(\xi_T Z\right).$$

If in addition $E\left(\xi_T Z^*\right) < +\infty$, then Z^* is the unique optimal solution almost surely. Thus, given a payoff X_T distributed with cdf F, the (a.s.) unique payoff Z_T also distributed with cdf F that minimizes the cost is $Z^* = F^{-1}\left(1 - F_\xi\left(\xi_T\right)\right).$

B.2 Probability of hitting the cap level

The probability of observing an index return over $[t, t + \Delta t]$ under the actual probability $\mathcal{P}$ greater than the cap return $c\%$ is:

$$\Pr\left(\frac{\mathrm{S_{t+\Delta t} - S_t}}{\mathrm{S_t}} > \mathrm{c\%}\right) = \Phi\left(\frac{-\ln(1+\mathrm{c}) + \left(\mu - \delta - \frac{\sigma^2}{2}\right)\Delta\mathrm{t}}{\sigma\sqrt{\Delta\mathrm{t}}}\right)$$

B.3 Overweighting Technique

Overweighting the tail of the distribution is equivalent to modifying the probability distribution or equivalently to changing the measure under which the investment outcomes are studied. Kahneman and Tversky (1992), Machina (1982), Brunnermeier et al. (2007), Ali (1977) work with discrete distributions. Barberis and Huang (2007) extend Tversky and Kahneman's (1992) model to continuous distributions.

In our case, we use Girsanov's theorem and the change of measure. There exists a probability measure Q equivalent to P such that $Z_t = W_t - \theta t$ is a Brownian motion under the new probability Q. Thus, the index price under the new probability Q can be written as

$$S_t = S_0 e^{\left(\mu - \delta + \theta\sigma - \frac{\sigma^2}{2}\right)t + \sigma Z_t}$$

The effect of the change of measure is only to change the drift of the geometric Brownian motion. If $\theta > 0$ then, the instantaneous return is higher by an additional instantaneous return of $\theta\sigma$ and the volatility is unchanged. Since only the drift changes, the probability of high outcomes is increased whereas low outcomes are less probable.

C Hypothetical Examples contained in the $JPL.G$ prospectus

This appendix corresponds to two pages extracted from the prospectus of the contract $JPL.G$ sold by JP Morgan Chase. We only reproduce the first two of the five examples contained in the prospectus.

Example 1

Hypothetical Examples of Amounts Payable at Maturity

In this example, we assume that the Index increases by 6% each quarter over the five-year term of the notes from a hypothetical initial Index closing level of 1150 to 3688 on the final quarterly valuation date. The quarterly capped Index return for each quarterly valuation period is 6% and the sum of those quarterly capped Index returns over the 20 quarterly valuation periods during the term of the notes is 100%. Thus, the cash payment at maturity for each $1,000 principal amount of notes will equal $2,200, consisting of the $1,000 principal amount plus an additional amount of $1,200 (120% times $1,000).

	2004			2005			2006		
Quarterly Valuation Period Ending:	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return
January				1292	6.0%	6.0%	1631	6.0%	6.0%
April				1370	6.0%	6.0%	1729	6.0%	6.0%
July	1150			1452	6.0%	6.0%	1833	6.0%	6.0%
October	1219	6.0%	6.0%	1539	6.0%	6.0%	1943	6.0%	6.0%

	2007			2008			2009		
	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return
January	2059	6.0%	6.0%	2600	6.0%	6.0%	3282	6.0%	6.0%
April	2183	6.0%	6.0%	2756	6.0%	6.0%	3479	6.0%	6.0%
July	2314	6.0%	6.0%	2921	6.0%	6.0%	3688	6.0%	6.0%
October	2453	6.0%	6.0%	3097	6.0%	6.0%			

Assumptions:	
Issue price per note (your investment):	$1,000
Sum of quarterly capped Index returns for the 20 quarterly valuation periods:	120.0%
Total return of the Index [(3688 – 1150) ÷ 1150] x 100:	220.7%
Calculations:	
Calculation of payment at maturity per $1,000 principal amount of the notes	
At maturity, you would receive a cash payment equal to the greater of:	
(x) $1,100	$1,100
and	
(y) $1,000 plus the additional amount $1,000 + ($1,000 x 120%)	$2,200
Investor receives $2,200 at maturity (120.00% total return on a hypothetical investment in the notes)	

As this example illustrates, due to a lack of compounding, investing in the notes is not the same as investing in the Index or the stocks underlying the Index, even where the Index consistently increases over the 20 quarterly valuation periods. Although the hypothetical increase of the Index for each quarterly valuation period in this example is 6%, which is not greater than the 6% cap on each quarterly capped Index return, the total return on the notes is 120%, whereas the total return on the Index for the same period is 220.7%. In this example, if you had invested $1,000 directly in the Index instead of the notes, you would have received $3,207 instead of the payment at maturity of $2,200 per note.

Example 2

In this example, we assume that the Index generally increases over the five-year term of the notes but has some negatively performing quarters. In this example, the sum of the quarterly capped Index returns over the 20 quarterly valuation periods during the term of the notes is 24.6%. Thus, the cash payment at maturity for each $1,000 principal amount of notes will equal $1,246, consisting of the $1,000 principal amount plus an additional amount of $246 (24.6% times $1,000).

	2004			2005			2006		
Quarterly Valuation Period Ending:	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return
January				1208	1.7%	1.7%	1527	13.4%	6.0%
April				1358	12.5%	6.0%	1558	2.0%	2.0%
July	1150			1264	–6.9%	–6.9%	1611	3.4%	3.4%
October	1187	3.2%	3.2%	1347	6.6%	6.0%	1587	–1.5%	–1.5%

	2007			2008			2009		
	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return	Hypothetical Index Closing Level	Quarterly Index Return	Capped Quarterly Index Return
January	1541	–2.9%	–2.9%	1868	0.2%	0.2%	1808	–1.0%	–1.0%
April	1754	13.8%	6.0%	1849	–1.0%	–1.0%	1845	2.0%	2.0%
July	1648	–6.0%	–6.0%	1888	2.1%	2.1%	1892	2.6%	2.6%
October	1863	13.0%	6.0%	1825	–3.4%	–3.4%			

Assumptions:

Issue price per note (your investment):	$1,000
Sum of quarterly capped Index returns for the 20 quarterly valuation periods:	24.6%
Total return of the Index [(1892 – 1150) ÷ 1150] x 100:	64.5%

Calculations:

Calculation of payment at maturity per $1,000 principal amount of the notes

At maturity, you would receive a cash payment equal to the greater of:

(x) $1,100	$1,100
and	
(y) $1,000 plus the additional amount	
$1,000 + ($1,000 x 24.6%)	$1,246

Investor receives $1,246 at maturity (24.66% total return on a hypothetical investment in the notes)

As this example illustrates, the 6% quarterly cap on any increases of the Index and the full exposure to any decreases of the Index can result in the return on the notes significantly trailing the return on the Index or the stocks underlying the Index. The 6% cap on positive quarterly capped Index returns can result in only a small participation in significant positive Index returns. Combining this limited upside participation with the full exposure to any and all negative quarterly capped Index returns, this example shows how a payment at maturity on the notes can be significantly less than an investment directly in the Index or in the stocks underlying the Index. In this example, if you had invested $1,000 directly in the Index instead of the notes, you would have received $1,645 instead of the payment at maturity of $1,246 per note.

References

Ali, M. (1977): "Probability and Utility Estimates for Racetrack Bettors," *The Journal of Political Economy*, 85(4), 803–815.

Barberis, N., M. Huang. (2008): "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices," *American Economic Review, Forthcoming*.

Brunnermeier, M.K., C. Gollier. J. Parker. (2007): "Beliefs in the Utility Function. Optimal Beliefs, Asset Prices, and the Preference for Skewed Returns," *American Economic Review*, 97(2), 159–165.

Camerer, C., H. Kunreuther. (1989): "Decision Processes for Low Probability Events: Policy Implications," *Journal of Policy Analysis and Management*, 8(4), 565–592.

Campbell, J. (2006): "Household Finance," *Presidential Address to the American Finance Association, Journal of Finance*, pp. 1553–1604.

Carlin, B. (2008): "Strategic Price Complexity in Retail Financial Markets," *Journal of Financial Economics (forthcoming)*.

Cox, J.C., H. Leland. (2000): "On Dynamic Investment Strategies," *Journal of Economic Dynamics and Control*, 24, 1859–1880.

Dybvig, P. (1988a): "Distributional Analysis of Portfolio Choice," *The Journal of Business*, 61, 369–393.

——— (1988b): "Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market," *Review of Financial Studies*, 1, 67–88.

Evans, R., R. Fahlenbrach. (2007): "Do Funds Need Governance? Evidence from Variable Annuity Mutual Fund Twins.," *Working Paper Dice Center WP 2007-17*.

Henderson, B.J., N. Pearson. (2007): "Patterns in the Payoffs of Structured Equity Derivatives," *Working Paper, AFA 2008 New Orleans Meetings*.

Jin, H., X. Zhou. (2008): "Behavioral Portfolio Selection in Continous Time," *Mathematical Finance*, 18(3), 385–426.

Kahneman, D., A. Tversky. (1979): "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263–291.

Koco, L. (2007): "Annuity Sales Slid 7% in 2006," *National Underwriter Life & Health*, April 2, 2007, 111(13), 8–9.

Palmer, B. (2006): "Equity-Indexed Annuities: Fundamental Concepts and Issues," *Working Paper*.

Rossetto, S., J. van Bommel. (2008): "Endless Leverage Certificates," *Working Paper presented at the ECB-CFS Research Network on "The Market for Retail Financial Services"*.

Shefrin, H. (2000): *Beyond Greed and Fear*. Harvard Business School Press.

Shefrin, H., M. Statman. (1993): "Behavioural Aspects of the Design and Marketing of Financial Products," *Financial Management*, 22, 123–134.

Stoimenov, P.A., S. Wilkens. (2005): "Are Structured Products 'Fairly' Priced? An analysis of the German Market for Equity-Linked Instruments," *Journal of Banking & Finance*, 29(12), 2971–2993.

Tversky, A., D. Kahneman. (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297–323.

Wilmott, P. (2002): "Cliquet Options and Volatility Models," *Wilmott Magazine*, pp. 78–83.