



*Charles A. Dice Center for
Research in Financial Economics*

Engineering Lemons

Petra Vokata,
The Ohio State University

Dice Center WP 2020-21
Fisher College of Business WP 2020-03-021

This paper can be downloaded without charge from: [http://
www.ssrn.com/abstract=3223427](http://www.ssrn.com/abstract=3223427)

An index to the working paper in the Fisher College of Business Working Paper Series is located at:
<http://www.ssrn.com/link/Fisher-College-of-Business.html>

Engineering lemons*

Petra Vokata[†]

Ohio State University

October 2020

Abstract

Recent complex financial products sold to households contradict the basic premise of canonical innovation theories: financial innovation benefits its adopters. In my 2006–2015 sample of over 28,000 yield enhancement products (YEP) the securities offer attractive yields but *negative* returns. The products lose money both ex ante and ex post due to their embedded fees: on average, YEPs charge 6–7% in annual fees and subsequently lose 6–7% relative to risk-adjusted benchmarks. Simple and cheap combinations of listed options often first-order dominate YEPs. Competition, disclosure, or learning do not eliminate this inferior financial innovation over my sample period.

JEL Classification: G4, G13, G14, G18

Keywords: Household finance; financial innovation; hidden costs; complexity; structured products.

*I am grateful for comments received from George Aragon, Utpal Bhattacharya, Amit Goyal, Tobin Hanspal, Harrison Hong, Matti Imanen, Matti Keloharju, Samuli Knüpfer, Juhani Linnainmaa, Brigitte Madrian, Adair Morse, Toby Moskowitz, Mikko Niemenmaa, Peter Nyberg, Terrance Odean, Neil Pearson, Vesa Puttonen, David Sraer, Antti Suhonen, Matti Suominen, Boris Vallée, Annette Vissing-Jørgensen, Tuomo Vuolteenaho, and seminar participants at the European Finance Association Meetings, Helsinki Finance Summit, Research in Behavioral Finance Conference, the University of California–Berkeley, the University of California, Los Angeles, Ohio State University, Arizona State University, Georgetown University, London School of Economics, Copenhagen Business School, BI Norwegian Business School, University of Virginia, Darden, Michigan State University, University of Illinois at Urbana-Champaign, and the Nordic Finance PhD Workshop. I thank the Academy of Finland, Nasdaq Nordic Foundation, the OP Group Research Foundation, the Foundation for Promoting the Finnish Securities Market, and the Foundation for Economic Education for generous financial support. I am grateful to the commercial data platform for providing the data, LexiFi for providing the pricing tool, and Aalto University for hospitality.

Declarations of interest: none.

[†]Ohio State University, 832 Fisher Hall 2100 Neil Avenue Columbus, OH 43210. E-mail: vokata.1@osu.edu. I completed a significant part of this project during a research visit at the University of California–Berkeley.

1 Introduction

Canonical theories of financial innovation share one key implication: innovation benefits adopters whose needs are not fully met by existing securities. Financial intermediaries engineer new products that investors demand, and as a result may improve social welfare.¹ A less sanguine view is that innovators do not intend to help investors but rather try to exploit them to extract rents. For example, intermediaries may innovate products to hide fees or encourage excessive risk-taking of confused investors. Although views like these are common among academics, policy-makers, and in society at large, careful empirical investigations of ex-ante inferior financial innovations have been scarce.² This paper aims to fill this gap and investigate cases of new financial products that are almost certainly not engineered to benefit their adopters.

Quantifying benefits of financial innovations is challenging because adopters' private values are unobserved. One may think that novel funds that ex post underperform existing funds do not add value, but funds differ on many dimensions that may be hard to quantify ex ante and that some investors may find valuable. Similarly, expensive (Henderson and Pearson, 2011) and complex (C  lerier and Vall  e, 2017) structured products may appear to be designed to exploit unsophisticated investors, but investor heterogeneity and unobserved preferences pose a challenge to this explanation.

I address these challenges with new data and a new empirical method. I study yield enhancement products (YEPs), which package high-coupon bonds with short positions in put options. Two aspects make YEPs an ideal setting for my empirical analysis. First, YEP payoffs at maturity are defined by a finite set of features that are all known ex ante and are independent of post-issuance actions of issuers or investors. I use novel data from the most comprehensive data provider that records all these features in a semi-standardized textual description. The data covers more than 28,000 YEPs issued

¹See, for example, Ross (1976) and Allen and Gale (1994). Even in cases in which financial innovation can make all agents worse off (Elul, 1995), the use of new securities is individually optimal for adopters. Tufano (2003) and Lerner and Tufano (2011) provide a review of the literature on financial innovation.

²For examples of this negative view, see Schoar (2012), Philippon (2016), and Volcker (2009). Zingales (2015) notes that "57% of readers of The Economist [...] disagree with the statement that "financial innovation boosts economic growth." Campbell (2006), in his presidential address, calls for further theoretical and empirical research of such "perverse" financial innovation. Frame, Wall, and White (2018), in their most recent survey of empirical work on financial innovation, mention only Henderson and Pearson (2011) as an example of the dark side of financial innovation.

over 2006–2015. Second, the payoffs of YEPs can be closely approximated by those of listed options and I can therefore use relative valuation to derive costs and possible benefits of YEPs compared with existing securities. Specifically, I use unambiguous relative payoff comparisons by constructing *counterfactual payoffs* that first-order *dominate* YEPs. In other words, I identify situations in which the YEP payoff is ex-ante dominated by existing securities across all its dimensions, in all states of the world, and for any investor.³

My central empirical findings are twofold. First, I find the embedded fees of YEPs are large enough for their ex-ante and ex-post returns to be *negative*. This finding implies that unless an investor values YEPs for hedging purposes, they are not beneficial, because any risk-averse investor would be better off investing in the risk-free asset. The second result, which is novel, is that the products are often first-order stochastically dominated by simple combinations of listed options. This finding rules out that dominated YEPs are beneficial even for hedging motives. Taken together, these results imply that under a minimum set of assumptions YEPs are not designed to benefit investors.

A prominent example is a product linked to JPMorgan Chase depicted in gray in Fig. 1. The product has a maturity of three months and offers a coupon of 11.5% per annum. At maturity, the product repays the principal in full if the price of JPMorgan Chase does not fall below 75% of its initial price at any time during the three months. Otherwise, the payoff is decreased by the decline in the stock price at maturity. Assuming a 6% annual market risk premium and given the product's estimated beta (0.5 at issuance), the expected gross return of the product is less than 1% over its term. Because the product charges more than 5% in embedded margin, its net-of-fee expected return is less than -4%.

Moreover, I find that investors could take on less risk and achieve a significantly higher coupon rate, 26.7% p.a., by creating a payoff dominating this YEP with listed options. This synthetic payoff is depicted in black and consists of a simple combination of two put options and risk-free lending. Importantly, to construct this payoff, investors need less skill than one needs to estimate fees of YEPs. They only need to understand

³In a different setting Egan (2019) studies reverse convertibles, a class of YEPs, which dominate other reverse convertibles. I focus on cases in which simple combinations of listed options dominate YEPs.

and combine payoffs of vanilla options, but they do not need to apply any option pricing techniques.

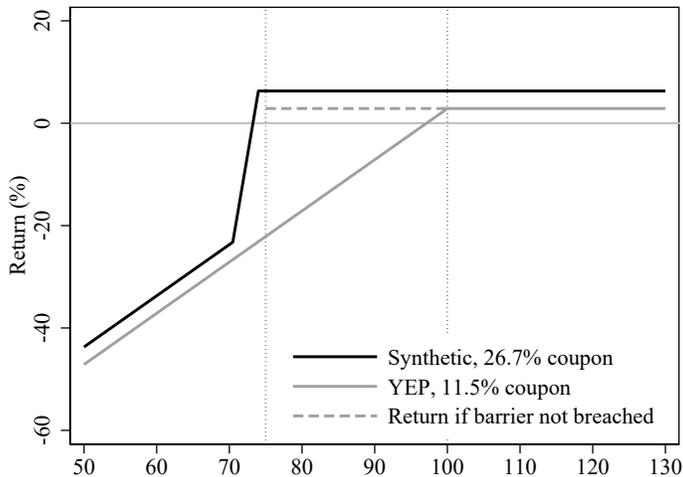


Fig. 1: Dominated example of yield enhancement product

The figure shows the total return diagram for YEP linked to JPMorgan Chase & Co. (CUSIP: 78008TYE8) and dominating synthetic security constructed using put options and lending. Prices of the underlying are normalized to 100 at issuance. Details for both securities are provided in Table 1.

I start my empirical analysis by providing new evidence on the costs and returns of YEPs. I first develop a precise translation algorithm to convert the textual descriptions of YEP payoffs into mathematical formulas. I then estimate the fair value of these payoff formulas using a local volatility diffusion model, which is a standard approach to value exotic payoffs. Depending on the YEP payoff, I use either a finite difference scheme, static replication, or Monte Carlo simulations to price path-dependent YEPs. Under conservative assumptions, on average YEPs embed 3.5% margin, which translates to a volume-weighted (equal weighted) average annual fee of 5.7% (7.2%). To put this number in perspective, YEPs are more than five times as expensive as equity mutual funds (Greenwood and Scharfstein, 2013).

Notably, these fees often exceed the gross expected return of YEPs. I quantify the expected returns from YEPs net of fees under various assumptions about the expected return on the underlying and find that, on average, YEPs are not expected to yield positive returns. One advantage of my sample size and period length is that I can evaluate if the poor YEP expected returns materialize into negative realized returns. To

do so, I evaluate the translated payoff formulas at maturity and any possible coupon payments and early terminations occurring over the life of the product. I find that both the total and annualized ex-post returns are, on average, negative. This poor ex-post performance is not limited to the financial crisis of 2007–08. For instance, the third of the products with the shortest maturities—and therefore with the shortest times to recoup their fees—earn negative average returns in eight out of the ten sample years.

To risk adjust the ex-post returns, I develop delta-hedged returns that take into account the non-linearities in YEP payoffs. The volume-weighted (equal-weighted) abnormal return, calculated as the difference between the product return and delta-hedged return, is -5.9% (-7.4%) and, therefore, comparable to the ex-ante fees.

The most common performance measure used in the industry is average annualized returns. I show this measure is ill suited to evaluate even the raw performance of the products, because conditional early terminations of the products—autocalls—lead to a significant upward bias in annualized returns. As an alternative measure that does not suffer from this bias and is easy to understand for retail investors, I develop time-series indexes of YEP gross and net performance. The indexes can also be regressed on an appropriate benchmark index to estimate alphas as alternative measures of abnormal returns. Using CBOE Put Index, which, like YEPs, involves selling of put options, I estimate YEP annualized alpha of -7.9% . In sum, I provide large sample evidence based on a number of alternative measures that YEPs charge 6–7% in embedded fees and subsequently lose 6–7% relative to risk-adjusted benchmarks.

I next examine the benefits of YEPs relative to listed options as in the example product shown in Fig. 1. For each product I create a synthetic YEP from up to two put options that statically approximate and statewise dominate the YEP payoff at maturity. I then determine YEP as dominated whenever the implied coupon of the synthetic YEP exceeds the YEP coupon. This simple method is transparent, model free, and depends on a minimum number of assumptions while allowing me to derive a sufficient measure that identifies cases of inferior financial innovations.

In a sample of 17,000 YEPs that can be approximated with up to two options, I find listed options dominate 30–45% of YEPs. Importantly, the construction of synthetic YEPs takes into account transaction costs in the form of bid-ask spreads, and additional

transaction costs such as commissions, tax motives, access to leverage, or minimum investment amounts do not explain their dominance.

I perform a battery of robustness checks that provide comfort about the validity of the data, payoff translation, and valuation model. I validate my calculated returns with returns reported by the platform and by an independent consulting firm. I also compare my estimates of fair values with estimates reported by issuers, fair values from a commercial pricing tool, and prices reported in TRACE for secondary repurchases by issuers. To give one example, the correlation between my estimates and TRACE prices is more than 99%, which confirms that the data, my payoff translation, and the valuation method are highly accurate.

The length of my sample period also allows me to assess several potential disciplining mechanisms that could eliminate YEPs with negative expected returns. In 2013, the U.S. Securities and Exchange Commission (SEC) required issuers to disclose their estimates of the fair product value to give investors a chance to understand the fees. I find the mandatory value disclosure is not associated with a significant decline in fees. The bad performance of the market during the 2008 financial crisis allowed investors to learn about the possibly neglected downside risks of YEPs. Although I find issuance volume of YEPs declined by about 40% in 2009, already in 2010 sales recovered to more than 90% of the 2008 level. Finally, competition among YEP issuers or from cheap online brokers of listed options had the potential to push down prices of YEPs. I find that over 2010–2015, the fees of YEPs significantly declined. However, even in 2015, the fees of YEPs were large enough for their expected returns to be negative.

Given the inferiority of YEPs, a natural question is why investors are buying them. Evidence from regulatory investigations shows banks target the products mainly at retail and often unsophisticated investors. I document three patterns consistent with issuers catering to investor biases. First, the high coupons of the products are saliently advertised often in the name of the product, whereas to quantify the possible losses, investors need to use option pricing techniques. Second, consistent with framing effects, the possible loss of principal at maturity is often not emphasized as a capital loss but instead reframed as the delivery of the underlying. Third, I show the underlying stocks are not chosen at random but overrepresent highly volatile stocks with a higher probability of

downside losses. I also find that brokers selling the products receive significantly higher commissions from selling YEPs than dominating listed options. All these patterns are consistent with banks engineering YEPs to charge high fees to unsophisticated investors who would be better off keeping their money in mattresses.

My study adds to four strands of research. First, my results add to the literature on financial innovation and provide novel evidence that financial innovation may not always benefit society (Zingales, 2015; Heidhues, Kőszegi, and Murooka, 2016; Gennaioli, Shleifer, and Vishny, 2012; Allen, 2012). Unlike in the contexts in which the suboptimal use of new financial products is what lowers household welfare (Bertrand and Morse, 2011; Li, Subrahmanyam, and Yang, 2018), my conclusions are independent of the actions of adopters and their unobserved preferences.

Second, my work relates to the literature on the cost of financial intermediation (Philippon, 2015; Hortaçsu and Syverson, 2004) and hidden fees (Gabaix and Laibson, 2006; Anagol and Kim, 2012; Duarte and Hastings, 2012). The evidence that YEPs yield negative returns suggests search costs or trust (Gennaioli, Shleifer, and Vishny, 2014) alone cannot explain their high costs and that behavioral explanations may play a role in investors' perception of product's intrinsic values. The results on mandated value disclosure add to the literature on regulation of consumer financial products (Campbell, Jackson, Madrian, and Tufano, 2011; Beshears, Choi, Laibson, and Madrian, 2011; Agarwal, Chomsisengphet, Mahoney, and Stroebel, 2015; Campbell, 2016) and show disclosure may not be enough to drive bad products out of the market.

Third, my paper relates to the growing literature on reaching for yield (Stein, 2013; Hanson and Stein, 2015; Di Maggio and Kacperczyk, 2017; Lian, Ma, and Wang, 2019). My results are consistent with behavioral explanations, such as framing and salience (Bordalo, Gennaioli, and Shleifer, 2016), playing an important role in reaching for yield by retail investors.

Fourth, my paper contributes to the literature on retail structured products (Henderson and Pearson, 2011; Bergstresser, 2008; Célérier and Vallée, 2017; Egan, 2019; Henderson, Pearson, and Wang, 2020b) and complex financial products (Carlin, 2009; Carlin, Kogan, and Lowery, 2013; Ghent, Torous, and Valkanov, 2019). My paper is the first to provide large sample evidence on the risk-adjusted performance of a class of

structured products; it identifies a new source of their performance manipulation, and develops a new method based on performance indexes to benchmark their performance.

The remainder of this paper is organized as follows. Section 2 describes my sample of YEPs. Section 3 presents the estimated costs and returns of YEPs. Section 4 describes the construction of synthetic and dominated YEPs. Section 5 shows that common performance measures used in the industry are biased and presents YEP performance indexes. Section 6 analyzes fair value disclosure. Section 7 discusses possible reasons why investors buy the products. Section 8 concludes.

2 Sample of yield enhancement products

YEPs—also categorized as income products—represent the largest category in terms of the number of products of retail structured notes offered in the U.S.⁴ Their issuance volume accounts for more than 40% of the volume of structured notes registered with the SEC. Banks market the products under different names, such as reverse convertible notes, income securities, yield optimization notes, equity-linked securities, and reverse exchangeable securities. In recent years, auto-callable securities—a class of YEPs that terminate early if the underlying rises above a predefined call price—have become more popular.

YEPs derive their return from the performance of the underlying asset or basket of assets. The most common underlying is a single stock or an equity index. A typical product has limited upside, determined by a fixed coupon rate. As the name suggests, this coupon rate—also called the headline rate—is higher than the prevailing interest rate. The higher yield is compensated by downside risk that is embedded in the product payoff through a short position in plain vanilla or exotic put options.

The products do not charge any ongoing fees but embed a margin of which part

⁴Other types of retail structured products include participation products and capital protected notes, studied, for example, by Calvet, C  lerier, Sodini, and Vall  e (2017). In the insurance market, structured equity-linked annuities represent more than one-third of insurers' liabilities (Kojen and Yogo, 2017). The term "structured (finance) products" is also used for structured finance vehicles that pool large numbers of economic assets and subsequently issue tranches against these collateral pools (Coval, Jurek, and Stafford, 2009). I use the term for securities that derive their payoff from a small number of underlying assets and a non-linear payoff formula.

goes to the distributing brokers in the form of a commission, and the remaining part compensates the issuers.

2.1 Data

My data on YEPs come from a commercial data provider that collects data on structured products issued all over the world. Following [C el erier and Vall e \(2017\)](#), I refer to it as "the platform." The platform is the most comprehensive source of data on retail structured products. At the time of data retrieval, it spanned over 50 countries, 16 years, and 11 million products.

I start my sample construction with 30,637 YEPs issued between January 2006 and September 2015 and linked to equity indices or stocks covered in OptionMetrics. I use this sample to train an algorithm to translate product textual descriptions to mathematical formulas. The final translated sample covers 28,383 products, which represents more than 90% of the original sample. Because mine is the first study to use the U.S. database of the platform and the first one to use its payoff descriptions for a performance analysis, the online appendix describes the data coverage and several validity checks I perform.

2.2 Payoff translation

The key feature of the data is a complete and concise description of the product payoff—in the form of short semi-structured text distilled from the long and complex disclosure in the prospectus. [Table 1](#) shows the description for the example product. The first sentence defines the underlying asset, followed by a description of the product cash flows before and at maturity.

The main challenge of the translation is the large variety of descriptions that reflect both the heterogeneity in product payoffs and the semi-structured nature of the descriptions—the same product payoff can be described in multiple ways. To reduce the dimensionality of the translation, I first detect and replace synonymous phrases in the corpus of the descriptions.⁵ I then validate and hand-translate the most common

⁵To pre-process the data, I remove upper cases and stop words and I strip the first sentence specifying

product payoffs. Panel B of Table 1 shows the final formula for the example product.

To evaluate the product values, I need to merge the translated formulas with data on underlying prices and valuation inputs from the Center for Research in Security Prices (CRSP) and OptionMetrics' IvyDB US database. To do that, I use the only identifier of the underlying asset available in the platform—the underlying name. For each name, I find the closest security name in OptionMetrics in terms of the Levenshtein distance (Levenshtein, 1966) and validate name pairs that are not perfect matches. I then merge CRSP with OptionMetrics using the CUSIP code.

I evaluate the payoffs using the dates of initial and final valuation recorded by the platform. The platform, however, does not report the dates on which the conditions for conditional coupons or early termination (knock-out) are determined. I either extract these observation dates directly from the prospectuses or I extract the coupon and knock-out frequency from the payoff description and extrapolate the dates from the initial valuation date.⁶

2.3 Payoff validation

I first verify that my formulas accurately describe product payoffs. To this end, I use three independent data sources. First, I compare my estimates of product fair values derived from the translated formulas with secondary market prices reported in TRACE. Next, I compare my ex-post returns with ex-post returns reported by the platform or by an independent consulting firm. The correlations between the estimates are 99%, 93%, and 96%, respectively, and I find the rare discrepancies in ex-post returns are driven by stock splits or reused tickers not factored in by the platform. Section A.3 of the online appendix provides further details on these checks.

the underlying and all numerical variables. I train a word2vec embedding model on n-grams of up to five words to detect synonymous phrases.

⁶I find that unlike in the bond market, in my sample, short/long last coupons are more common than short/long first coupons.

2.4 Descriptive statistics

Table 2 presents summary statistics of the translated sample. Panel A summarizes product characteristics. The average headline rate is 12%—an order of magnitude higher than the prevailing interest rate. The overnight indexed swap (OIS) rate with matching maturity averages only 1.4%. The products have short maturities. Their average maximum term—if they do not terminate early—is one year. Panel B reports the average underlying factor loadings from the Fama and French (2015) five-factor model. The underlyings are typically highly volatile stocks selected systematically to support high headline rates and moderate downside protection.⁷ Their average beta is 1.6—a value common for the top beta decile of the U.S. stocks.⁸

3 Costs and returns of YEPs

In this section, I estimate YEP costs and returns. I first discuss the estimated ex-ante measures—embedded margins, fees, and implied expected returns—and then turn to the ex-post evidence on YEP performance.

3.1 Valuation model

I estimate the values of YEPs using a local volatility diffusion model, which is a standard industry approach to value exotic payoffs and has been used by (C  lerier and Vall  e, 2017). The price of the underlying asset, S_t , follows the diffusion:

$$dS_t = r_t S_t dt + \sigma(t, S_t) S_t dW_t + dA_t, \quad (1)$$

where r_t is the interest rate, $\sigma(t, S_t)$ is the volatility surface as a function of time t and spot level S_t , W_t is a Brownian motion, and

$$A_t = - \sum_{i=1}^{n_D} c_i 1_{t \leq t_i}, \quad (2)$$

⁷Section 7.2 provides more evidence on the underlying selection.

⁸See Table A.4 in the online appendix for the 40 most frequent underlyings and Table A.3 for the 15 largest issuers.

where $0 < t_1 < t_2 < t_{n_D}$ are ex-dividend dates and c_i are dividend amounts.

To calculate local volatility, $\sigma(t, S_t)$, with Dupire's formula (Dupire, 1994), I use a discrete set of implied volatilities from volatility surface of OptionMetrics and non-parametric arbitrage-free interpolation following Andreasen and Huge (2011). OIS rates to build the yield curve are from Bloomberg (Hull and White, 2013). Unless a dividend payment date is already declared, I project the ex-dividend dates by extrapolating from the previous 12 months of history. For indexes, I assume a constant dividend yield equal to the dividend yield reported in OptionMetrics.

I then follow industry practice and use the local volatility model described above with an appropriate valuation method for each payoff type. That is, I use a finite difference scheme for products with embedded vanilla or barrier options, static replications for digital options, and Monte Carlo simulations for autocallable products. Section B of the online appendix contains further details.

Because the products expose investors to the default risk of the issuer, their values should be adjusted for credit risk. A common proxy for the issuer credit risk is the CDS spread, which is not available for a quarter of the products in my sample. In addition, such credit value adjustment is too large when issuers' default is negatively correlated with YEP payoffs, which is likely the case for the products in my sample. For these reasons, I estimate the value of the products without the credit value adjustment, but discuss the size of the impact of the adjustment for products with available CDS spread. CDS data are from (in order of priority) CMA Datavision, Thomson Reuters, and Bloomberg.

3.2 Validation

Before describing the results, I first perform three independent checks to verify the accuracy of this computationally intensive valuation. First, I compare my estimates with product fair values from a commercial pricing tool used by product issuers as well as by academics (C  lerier and Vall  e, 2017). This comparison allows me to keep the product data, valuation inputs, and the broad valuation approach fixed with differences stemming only from minor differences in implementation. I confirm that the estimates are highly positively correlated ($\rho = 95\%$, $N = 21,390$) with an economically unimportant mean

difference of 2.9 basis points.

As a second check, I compare my estimates with the estimates disclosed by issuers following the 2013 SEC requirement. I find issuer estimates are also highly correlated with my estimates ($\rho = 87\%$ for annual margins and $\rho = 61\%$ for fair values, $N = 6,372$), despite the fact that they can be based on different valuation inputs or approaches. Finally, the high correlation between my estimates and secondary prices reported in TRACE ($\rho = 99\%$, $N = 24,746$) serves as a third check of my valuation approach. Section B.1 of the online appendix describes these checks in more detail.

3.3 Margins and embedded fees

Table 3, Panel A, reports the estimated product margins at issuance. I define the margin as

$$\text{margin} = \frac{\text{price} - \text{fair value}}{\text{price}}. \quad (3)$$

The average margin before adjusting for the credit risk is 3.6%; that is, a product sold for \$1,000 is on average worth only \$964.

In the second and third columns, I estimate the margins only for the subsample of products with available CDS spread data. The difference between the unadjusted ($r = r_f$) and the credit-risk-adjusted ($r = r_f + CDS$) margins is, on average, only about 0.14 percentage points. Because the effect of credit value adjustment is small and the CDS data are not available for a significant fraction of the products, I focus on the unadjusted values in the rest of the paper. These fair values, therefore, represent lower bound estimates, and the expected returns derived from them represent upper bound estimates.⁹

In Panel B, I convert the margins into annual embedded fees. Previous literature sometimes reports the annualized margin (C  lerier and Vall  e, 2017; Calvet, C  lerier, Sodini, and Vall  e, 2017), but this measure does not take into account early termination of products, which makes them more expensive on an annual basis. I calculate annual

⁹If an investor decides to sell the product before maturity, any markdown charged by the issuers in the secondary prices constitutes another source of embedded costs. Section C.1 in the online appendix estimates that the markdowns average 2% of the secondary fair value.

fees as

$$\text{fee} = \frac{\text{margin}}{E(\text{maturity})}, \quad (4)$$

where $E(\text{maturity})$ is the expected maturity of the product calculated using risk-neutral probabilities of termination. On (volume-weighted) average, investors pay (5.7%) 7.2% in annual fees. As a point of reference, YEPs are more than five times as expensive as equity mutual funds (Greenwood and Scharfstein, 2013).

I find the product expected term is the most important determinant of its fees. Although both the margin and commission increase with the product term, they do not scale linearly with the product term but include a fixed component. As a result, the shorter the product term, the more expensive the product is on an annual basis. Columns 2 – 4 in Panel B show that for short-term products with an expected term up to four months, the average annual fee is as high as 11 – 12%, whereas longer-term products with an expected term above eight months have average fees below 4%.¹⁰

3.4 Expected returns

I now test how the fees reported in the previous section affect the net-of-fee expected returns of YEPs and show that under various measures of the expected return on the underlying, the majority of the products in my sample have a negative expected return. To this end, I extend the pricing model described in Section 3.1 and calculate the expected undiscounted product payoffs under the objective ("real-world") expected return on the underlying, μ . I update the underlying diffusion as follows:

$$dS_t = \mu_t S_t dt + \sigma(t, S_t) S_t dW_t + dA_t. \quad (5)$$

The estimated product payoff expressed as a percentage of the issue price minus one equals the product expected return net of fees. In my baseline specification, I estimate

¹⁰The margins I find are in line with the YEP margins in Henderson and Pearson (2011) and Egan (2019) in the U.S. market. To be clear, the estimated fees are specific to the U.S. YEPs and may not extend to other retail structured product markets. For example, given the important relationship between product term and fees, capital protected products, which tend to have longer maturities, are likely significantly cheaper than YEPs. Previous literature also shows substantial variation in structured product markets across countries (C el erier and Vall ee, 2017; Baule, 2011; Szymanowska, Horst, and Veld, 2009; Burth, Kraus, and Wohlwend, 2003).

the expected return on the underlying asset using the CAPM β estimated over the past 60 months and a 6% p.a. market risk premium. The first column of Table 4 reports the results and shows the sample average, the volume-weighted average, and the median are all negative, ranging from a total return of -0.5% to -1.5% . The estimated returns are higher with the market premium equal to 8% p.a. (second column) or to the value-weighted CRSP average (third column), but even in these specifications, the median expected return is negative and the average is not significantly different from zero.

In the previous analysis, I assume the market risk premium is constant or equal to the historical average. Martin (2017) shows his measure of expected market return (SVIX) derived from option prices implies a large time-series variation in expected market return, which exceeds 20% in the peak months of 2008. I estimate product expected returns using SVIX for a sample of products issued before February 2012 (due to data availability). Column 4 confirms the average expected return is negative even under SVIX. Moreover, the individual averages of expected returns in nine out of the ten years of my sample estimated with a 6% p.a. market risk premium are negative as well. Therefore, my results are not driven by the unusual market conditions in the fall of 2008.

In the last column of Table 4, I use the Fama and French (2015) five-factor model instead of the single-factor model. A large empirical literature dating back to Black (1972) shows underperformance of high-beta stocks relative to the CAPM predictions. Given that high-beta stocks are overrepresented in the sample of underlying equities, the single-factor model likely overestimates the expected returns of the underlying stocks. This is consistent with the expected returns estimated using the five-factor model, which are significantly lower than in the previous specifications. The volume-weighted average expected total return in this specification is -1.73% , and for three-quarters of YEPs, I estimate a negative expected return.

In Panel B of Table 4, I report the estimated expected returns on an annual basis calculated as the internal rate of return that makes the discounted expected payoffs equal to the issue price. The expected annualized returns are negative in all specifications.

3.5 Ex-post performance

The evidence presented so far focuses on ex-ante costs and returns. In the rest of this section, I provide the first comprehensive evidence of YEPs' ex-post performance.¹¹

Due to the paucity of secondary market trading of YEPs, time-series data on their returns are scant. For this reason, I start by analyzing YEP returns at maturity. Although the platform records realized returns provided by the issuers or calculated by the platform analysts, its coverage is only about 50% of my sample. To get a comprehensive record of YEP performance, I calculate their returns using the translated payoff formulas combined with the ex-post prices of the underlying assets correctly adjusted for stock splits and stock dividends. The realized return is defined as the sum of all coupons and the payoff at maturity divided by the issue price minus one. This approach does not take into account the possibility of issuer default, which is rare but happened, e.g., in case of products issued by Lehman Brothers.

To correctly risk adjust YEP's non-linear returns, I calculate benchmark returns from delta-equivalent daily adjusted portfolios of the underlying and the risk-free bond (Broadie, Chernov, and Johannes, 2009). Daily delta,

$$\Delta_{i,t}(S_{i,t}, \sigma(t, S_{i,t}), r_t, A_{i,t}) = \frac{\partial \text{fair value}_{i,t}}{\partial S_{i,t}}, \quad (6)$$

comes from the valuation model described in Section 3.1. I cap the absolute product delta at two to avoid extreme positions. For each product i and trading day t , I calculate the daily benchmark return, $r_{i,t}^b$, as

$$r_{i,t}^b = r_t + \Delta_{i,t}(r_{S_{i,t}} - r_t), \quad (7)$$

where $r_{S_{i,t}}$ is the return on the underlying stock on day t . The benchmark return at maturity is then the cumulative return on this delta-equivalent replicating portfolio.

Table 5 presents the results. Over the sample period, investors in YEPs lost money on average. The volume-weighted average total return is -3.98% , or -3.46% annually.

¹¹Previous large-sample studies, e.g., Deng, Dulaney, Husson, McCann, and Yan (2015), analyze ex-post performance but do not quantify risk-adjusted performance that takes into account non-linearity in YEP payoffs.

Over a quarter of the products paid back less than the invested capital. Products issued in the years 2007 and 2008 earn the lowest average returns, but the negative returns are not confined to the crash period. Investors lost money on average even in the years 2011, 2014, and 2015, when the market earned positive returns. On a risk-adjusted basis, the total abnormal returns (defined as the difference between the product return and benchmark return) are -3.25% , or -5.93% annually, and therefore comparable to the ex-ante estimated margins and fees.

4 Synthetic YEPs and dominated products

The previous section shows that YEPs charge fees large enough for their expected and realized returns to be negative. Such products could still be desirable for investors if they are intended for hedging or speculative motives. Consistent with such motives, the products are often advertised as providing access to exotic payoffs. Although the exotic payoffs of YEPs are not easily available to retail investors, in this section, I show payoffs that first-order dominate YEPs are often easily available through listed options.

4.1 Construction of synthetic YEPs

To this end, I construct a synthetic counterpart to each YEP that can be statically approximated with up to two positions in put options and lending, as in the example product in Fig. 1. The construction of these synthetic YEPs varies depending on the exact YEP payoff but follows a simple algorithm. First, from all listed put options with the same underlying available on the initial YEP pricing date, find the ones with the closest expiry date before the YEP maturity. Next, find two options with the closest strike prices below the YEP barrier level. Finally, buy and sell the quantities in these two options that will most closely approximate from above the YEP payoff at maturity excluding any coupon payments, and lend the nominal value of YEP for the risk-free rate. Fig. D.3 in the online appendix and Panel C of Table 1 provide an example of the construction of synthetic YEP.

As price inputs, I use bid and ask option prices and the OIS rate as a proxy for the risk-free rate. The sum of the risk-free rate plus the net income from selling and buying

the put options annualized using the option expiry date and 30/360 day count convention then gives the coupon rate of this synthetic YEP. Its most important characteristic is that as long as the synthetic coupon rate is higher than the original YEP coupon rate, the synthetic product statewise dominates the YEP.

The construction of the synthetic counterparts is intentionally simple so that it could be easily implemented by any retail investor who understands the payoffs of YEPs. It is not intended to best approximate the payoffs of the products, which could be better achieved with more sophisticated quasi-static replications requiring more skill. Note this approximation does not require any knowledge of option pricing. All it requires is an understanding of the payoffs of a short and long position in a put option and their combination. It also does not require any additional margin, because the maximum possible loss is fully cash secured by the nominal value of the product invested at the risk-free rate.

4.2 Evidence on dominated YEPs

Fig. 2, Panel A, compares the coupon rates of YEPs with their synthetic counterparts. The coupon rate of YEPs averages 12.97% p.a., whereas the synthetic YEPs average 11.62%. This comparison shows that although synthetic YEPs offer in many cases significantly better downside protection, their coupon rates are only 1.35 percentage points lower.

The downside protection of synthetic YEPs is significantly better for three reasons. First, the synthetic YEPs do not include path-dependent conditions, which make the protection weaker. For example, for the barrier product in Fig. 1, YEP exposes investors to losses of the principal if the underlying drops below the barrier on any day during the life of the product, whereas the synthetic YEP exposes investors to principal losses only if the underlying is below the barrier at maturity. Second, if listed options do not span the barrier level of YEPs, the level of downside protection of YEPs is by construction lower because the algorithm takes the first option with strike price below the barrier level. On average, the synthetic YEPs protect the principal up to a 29.2% fall in the underlying, whereas the corresponding protection of YEPs is only up to 25.6%. Finally, because the sharp decline below barrier is not perfectly vertical, synthetic YEPs offer

better protection even below the barrier.

Given the lower downside risk of synthetic YEPs, the resulting small difference between synthetic and YEP coupon rates casts some doubt on the view that YEPs are access products or that their payoffs allow investors better reach for yield in a low-interest rate environment (C  lerier and Vall  e, 2017). In fact, for nearly half of the sample, investors would have achieved higher yield with the synthetic approximations of YEPs.

In Panel B of Fig. 2, I define the excess synthetic coupon rate as the difference between the synthetic and YEP coupon rate. Therefore, when the excess synthetic rate is positive, the synthetic payoff statewise dominates the original YEP. I call these products dominated YEPs. The figure shows 43% of the products in the sample are dominated. Column 1 of Table 6 shows that on a volume-weighted basis, the fraction of dominated products is even slightly higher—45%.

In other words, for nearly half of the products, investors would be better off in all states of the world by investing in a simple combination of put options and risk-free rate rather in the complex YEP.¹² In these cases, investors would get better downside protection and a higher coupon rate and at the same time enjoy higher liquidity, better transparency, greater customization, and lower risk of market manipulation through pre-trade hedging (Henderson, Pearson, and Wang, 2020b) of listed options. I show other differences, such as between American- and European-style exercise, minimum investment amounts, tax reasons, or leverage, do not make YEPs preferable to their synthetic counterparts either. Therefore, rationalizing dominated YEPs with any theory of financial innovation intended to reduce transaction costs or to improve spanning is hard.

One difference between YEPs and standardized options is that YEPs have European-style exercise whereas stock options are American. Exercise of deep-in-the-money American put options can be optimal to earn interest. Given the low interest rates during my sample period, the fact that most of the options at issuance are deep out of the

¹²Table D.4 in the online appendix shows the probability of being dominated as well as the excess synthetic coupon rate are positively related to brokers' commissions and issuers' margins. In addition, the probability of being dominated is the highest for product payoffs with embedded vanilla options, which are most easily approximated with listed options, and the lowest for products with barrier options, which cannot be statically replicated with vanilla listed options.

money, and that investors often fail to optimally exercise put options (Barraclough and Whaley, 2012), this concern is not important. Nevertheless, I examine the impact of early exercise on the synthetic payoffs.

Investors in the synthetic YEP facing an early assignment of the short put position can execute the long position and finance the residual assigned quantity with the nominal value deposited at the risk-free rate. The final synthetic payoff will differ only in the amount of income from lending at the risk-free rate over the remaining time to maturity. Column 2 of Table 6 shows that under a significantly more conservative assumption of excluding all lending income (i.e., even for the majority of cases in which the embedded options are never deep in the money and for the whole duration of the synthetic YEP), about one third of YEPs remain dominated by synthetic YEPs.

As another robustness check, I examine the role of different maturity between YEPs and synthetic YEPs. To annualize the coupon rate of synthetic YEPs, I use the maturity of its option components, which is by construction shorter than the YEP maturity, whenever listed options do not perfectly span YEP maturities. I believe this approach is reasonable given that YEP investors do not appear to have a strong preference for the exact maturity of YEPs and often invest in YEPs that can terminate early. Nevertheless, I next evaluate the impact of the differences in maturities. Fig. D.5 in the online appendix plots the difference between the YEP maturity date and the closest expiry date of listed options. Out of the 17,000 YEPs with available synthetic counterparts, more than half (8,915) have available listed options that expire up to 40 days before YEP maturity. This evidence goes against the idea that the main purpose of YEPs is to span maturities unspanned by listed options. Column 3 of Table 6 shows that in these 8,915 products, 28–30% of YEPs remain dominated even if I annualize the synthetic coupon rate using the longer YEP maturity. Note that in this specification, the downside risk exposure of YEPs compared to the synthetic YEPs is even larger, because YEPs expose investors to downside losses over a longer period.

One may also wonder if differences between the minimum invested amounts of YEPs and synthetic YEPs play a role. Standardized option contracts are for 100 shares of the underlying stock, which implies an average minimum position in a synthetic YEP of \$6,423. The minimum investment amount of YEPs is usually \$1,000, but the typical

invested amounts are significantly higher. For example, the average investment by 8,700 investors in UBS reverse convertibles at the center of the SEC fine for unsuitable sales was \$63,000.¹³ The median invested amount implied from my sample of secondary TRACE prices is \$25,000. Moreover, Columns 5 and 6 of Table D.4 show that dominated YEPs have a slightly higher issuance volume, although the difference is not statistically significant. These patterns mean dominated YEPs are unlikely improve the access of investors by requiring lower investment amounts than synthetic positions in listed options.

The high invested amounts in YEPs also imply commissions and per-contract fees of listed options will have a small impact on the profitability of synthetic YEPs. YEPs also do not offer any tax benefits or better access to leverage.

5 Biased performance measures and YEP indexes

Section 3.5 uses delta-hedged returns to show that YEPs significantly underperform risk-adjusted benchmarks. This result is at odds with the frequent claims by YEP promoters that the products can provide attractive or superior performance. The evidence provided to investors is often based on average or median annualized historical return. These measures are only sporadically accompanied by precise risk-adjusted benchmarks and have been used by academics.¹⁴ This lack of precise risk-adjustment that takes into account the non-linearity in structured product payoffs as well as their performance-contingent early maturity make these measures misleading.

If investors rely on them, issuers of the products have an incentive to create products that fare well on these measures—either through lower fees, which increase expected returns, or by engineering security designs that spuriously outperform. A simple design feature that manipulates the measures is performance-contingent early termination, that is, an autocall feature.

Consider the following simple autocallable product: with a 50% probability, the product terminates in three months and returns 110% of the issue price. Otherwise, the product terminates in 12 months and returns 80% of the issue price. The average total

¹³See SEC order available at <https://www.sec.gov/litigation/admin/2016/34-78958.pdf>.

¹⁴See, e.g., FTAdviser article available at <https://www.ftadviser.com/investments/2019/10/31/how-structured-products-can-perform-in-a-shaky-economy/> and C  lerier and Vall  e (2017).

realized returns of the product converge to -5% , but the average annualized realized returns converge to 10% , despite the fact that the product has a negative expected return. Clearly, in this case, average annualized returns are not a reasonable measure of the product's performance.¹⁵

The magnitude of the bias introduced by early terminations is substantial. For example, in my sample of 7,586 autocallable products, the average total return is -0.84% , whereas the average annualized return is 5.90% . When term-weighting annualized returns, which partially corrects for the bias, the average annualized return is -1.81% .

This manipulation mechanism can be thought of as a dynamic manipulation strategy intended to induce estimation error in performance measures (Goetzmann, Ingersoll, Spiegel, and Welch, 2007). It will manifest itself in an upward bias of any risk-reward ratio (e.g., Sharpe ratio) based on annualized returns. Even absent this bias, such ratios are ill-suited in the context of structured products with non-linear payoffs. Benchmark-based measures, such as delta-hedged returns, are immune to this bias. The average abnormal annualized return based on delta-hedged returns in my sample of autocallables is -6.96% . In the next section, I develop an alternative manipulation-proof performance measure based on daily indexes of YEP performance.

5.1 YEP performance indexes

Time-series performance charts are extensively used when marketing and analyzing historical performance of retail financial products, and investors could therefore find them easier to interpret than delta-hedged returns. In addition, return series can be used to estimate widely used benchmark measures such as Jensen's alpha. When this benchmarking is based on broad stock market indexes, it also allows us to assess the impact of underlying selection or price manipulation (Henderson, Pearson, and Wang, 2020b), which delta-hedged returns do not capture.

An obvious challenge to creating an aggregate index of YEP performance is that

¹⁵Sample selection can introduce another bias, a type of survivorship bias, when analyzing products that already matured and therefore have a positive performance and not accounting for products that did not mature yet and will likely have a negative performance. This concern is not important in my analysis, because I observe ex-post returns for 97% of the products, but it is a common problem in performance studies focusing on the most recent period.

time-series data on their performance are not available, because the products are typically not listed and only sporadically traded. To address this challenge, I leverage my valuation model and estimate the fair value of the products on each day until their maturity. I then use these fair values to create two YEP performance indexes: net index, which takes into account the product embedded fees, and gross index, which excludes the effect of fees.

To construct the indexes, I follow the MSCI fixed-income index methodology¹⁶ with daily rebalancing and daily reinvested coupons. The indexes are volume weighted and cover 27,578 products with complete fair value history. The net index covers the products already since the initial issue date and therefore captures the drop in their value from issue price to fair price, whereas the gross index skips the issue day and therefore does not take into account initial overpricing.

Panel A of Table 7 presents the summary statistics of daily returns of both indexes as well as of the Cboe S&P 500 PutWrite (PUT) Index, which is investable through ETFs and a reasonable benchmark for YEPs. The PUT index measures performance of a strategy similar to YEPs that sells a sequence of one-month, at-the-money, S&P 500 Index put options and reinvests the proceeds at the T-bill rate.¹⁷ Both YEPs and the PUT index expose investors to potentially large downside losses and a limited upside, and they have the same degree of negative daily skewness. The YEP indexes are mostly linked to single-name stocks and are therefore less diversified, which is reflected in their higher standard deviation. On an annualized basis, the PUT index returned 6.36% in excess of the risk-free rate over the period from May 2006 to December 2016, while the gross and net YEP indexes returned only 3.93% and -0.90% , respectively.¹⁸

Fig. 3 plots the log index growth for all three indexes. Visually, the PUT and gross YEP strategies earn positive returns in quiet times but incur large negative returns around events such as the 2008 financial crisis, 2011 European debt crisis, or 2015–2016

¹⁶The methodology from November 2019 is available at https://www.msci.com/eqb/methodology/meth_docs/MSCI_FI_Index_Calculation_Methodology.pdf.

¹⁷For more details, see <http://www.cboe.com/products/strategy-benchmark-indexes/putwrite-indexes/cboe-s-p-500-putwrite-index-put>.

¹⁸The return of the net index is higher than average total or annualized product returns, because the index weighs longer-term products more heavily and gives equal weight to all days in the sample period. Fig. D.4 in the online appendix contrasts performance of the manipulation-proof index with the biased average annualized returns by product type.

stock market sell-off. The net YEP index does not perform well even in relatively quiet times.

In Panel B of Table 7, I use the indexes to estimate YEP abnormal returns. To this end, I regress the daily excess YEP returns on the PUT index excess returns and the size, book-to-market, and momentum factors. Columns 1 and 2 report regressions of the YEP Gross index, and columns 3 and 4 report regressions of the YEP Net index. The constants are annualized alphas.

The gross index has a negative but insignificant alpha of -3.12% . The underlying of YEPs could underperform the market for a number of reasons. First, [Henderson, Pearson, and Wang \(2020b\)](#) show issuers manipulate underlying prices on the initial issue date through pre-trade hedging, which should lower the subsequent underlying returns. Second, the issuers select underlyings with expensive out-of-the-money put options, which tend to underperform ([An, Ang, Bali, and Cakici, 2014](#)). Third, the products are linked to a selected number of underlyings that could perform worse than the market by chance ([Bessembinder, 2018](#)) or because of investor sentiment ([Henderson, Pearson, and Wang, 2020a](#)).

The alpha of the net index, however, shows most of the underperformance of YEPs comes from their fees. The net index has an annualized alpha of -7.91% , and the difference to the alpha of the gross index, 4.8% , is close to the ex-ante fees. In sum, independent of whether retail investors rely on benchmark-adjusted or raw returns ([Ben-David, Li, Rossi, and Song, 2019](#)), the net index is useful to communicate the poor track record of YEPs relative to investable alternatives.

6 Mandatory fair value disclosure

To give investors a chance to understand the undisclosed costs of YEPs, in 2013 the SEC required issuers to disclose their estimates of the fair product value. To the extent that the disclosure successfully unshrouds the hidden costs of YEPs to uninformed investors, it could drive expensive products out of the market. It could cause a decline in embedded fees, margins, or issuance volume. In this section, I examine these hypotheses and the accuracy of the disclosed values.

The intention to introduce fair value disclosure was first mentioned in a comment letter sent by the SEC in April 2012, but it was not introduced until after a second letter with more detailed instructions sent in February 2013.¹⁹ The SEC asked issuers to disclose their estimate of the fair value of a product, which can be derived using the issuer's internal funding rate, mid-market inputs of comparable derivatives (with possible adjustments), and proprietary pricing models. Issuers implemented the disclosure gradually over the second and third quarters of 2013.

Panel A of Table 8 shows the summary statistics of the disclosed issuer estimates for 6,372 products and compares them with the estimates from my valuation model. I find both estimates are very close to each other. The volume-weighted average issuer estimate is 96.92%, whereas my estimated fair value is 96.78%. The mean difference between the estimates, 14 basis points, is economically small and represents less than 5% of the estimated margin. I hence conclude the disclosed issuer values provide reasonable estimates of the product fair values.

I next examine the evolution of fees over the sample period in Fig. 4. The fees were lower at the beginning of my sample period, elevated around the period of the financial crisis of 2007, and started to decline again from 2011 onward. The evolution of average underlying implied volatility (measured at 1-year maturity and delta of -20) broadly matches the same time-series pattern, suggesting that market conditions may influence the market for YEPs. The figure also suggests that, at a first glance, the mandatory fair value disclosure did not have a dramatic impact on product fees. The long-term trend in declining fees started three years prior to the mandatory disclosure, and no discernible dramatic change occurs in 2013.

Table 8, Panel B, provides a more detailed analysis of the effects of the disclosure. In columns 1, 3, and 5, I test for the difference in fees, margins, and volume, respectively, in the four quarters before the introduction of the disclosure (2012q1–2012q4) and four quarters after issuers started to disclose their estimates (2013q3–2014q2). Although these differences do not have a causal interpretation, they are indicative of possible changes in market outcomes shortly after the disclosure. I find both the fees and margins

¹⁹Fig. D.1 in the online appendix presents an example of a pricing supplement that includes the disclosed issuer estimated value.

declined following the disclosure, but neither of the differences is statistically significant. The decline in fees of 36 basis points corresponds to 5.6% of the pre-disclosure average. The coefficient on volume is significant, but it has a positive sign. The average volume increased by \$0.4 million following the disclosure.

These pre-post differences may mask important heterogeneity in issuers. In particular, one may expect the impact of the disclosure to be greater for more expensive issuers. To test this conjecture, I estimate the following difference-in-differences specification:

$$y_{ikt} = \beta \text{High-Fee}_k \times \text{Post}_t + \text{Post}_t + \lambda_k + \epsilon_{ikt}, \quad (8)$$

where y_{ikt} is either the fee, margin, or volume sold of product i issued in quarter t by issuer k , High-Fee_k is an indicator equal to 1 for issuers with above-median fees in year 2012, Post_t is an indicator equal to 1 from 2013q2 onward, and λ_k stands for issuer fixed effects.

The coefficient β identifies the impact of the disclosure on high-fee issuers under the assumption that they would have maintained parallel trends with low-fee issuers in the absence of the disclosure. Consistent with this assumption, Fig. 4, Panel B, shows both groups have parallel trends before the introduction of the disclosure.

Columns 2, 4, and 6 report the results of this estimation. I find again that the effects on high-fee issuers are not statistically significant. Moreover, the coefficient on fees is positive, implying high-fee issuers charge even relatively higher fees after the disclosure. I conclude the disclosure did not have a significant effects on the costs of the products.

7 Why do investors buy YEPs

Given their negative expected returns, why would anyone buy overpriced YEPs? Although a definitive answer is beyond the scope of this paper, I am able to rule out rational explanations for YEP demand. Any rational investor demanding YEP payoffs would be better off investing in dominating portfolios of listed options and would need less skill to trade vanilla options than one needs to estimate fees of exotic YEPs.

I next turn to a few alternative behavioral explanations. I first focus on the char-

acteristics of the investors and then analyze the framing and marketing of YEPs.

7.1 Investor characteristics

Banks target the products mainly at retail investors (Aguilar, 2015). Products sold to accredited investors—with incomes above \$200,000 or net worth over \$1 million—are exempt from the SEC registration and therefore are not in my sample of YEPs. By FINRA rule 2111, broker-dealers can only sell products that are suitable for a customer based on the customer's investment profile. The recommendation of YEP is considered suitable only if one can reasonably believe the investor is capable of evaluating its risks based on her knowledge and experience. The internal suitability guidelines of certain broker-dealers expect investors to have at least two years of investment experience, a \$100,000 income, \$100,000 in liquid assets, and \$250,000 net worth.²⁰ Therefore, investors for whom YEPs are suitable, should be also easily approved for option trading.

Regulatory investigations show these suitability requirements are frequently violated. For example, the SEC and FINRA found that some investors do not understand the terms of the products and some broker-dealers are aggressively marketing the products to elderly, non-English-speaking investors and to investors with conservative investment objectives, modest income or wealth, and little investing experience.²¹ This evidence is consistent with the explanation that inferior YEPs are aggressively sold to unsophisticated households who are unlikely to understand the high costs and low returns of YEPs.

7.2 Evidence on catering to biases

A number of recent studies show an important role of framing, teaser rates, and advertising content in shaping consumer financial decisions (Ausubel, 1991; Bertrand, Karlan,

²⁰These criteria were described as "must" until 2009 and as "should" thereafter by RBC; http://www.finra.org/sites/default/files/fda_documents/2010022918701_FDA_JM992805.pdf.

²¹See SEC reports available at <https://www.sec.gov/news/studies/2011/ssp-study.pdf> and <https://www.sec.gov/about/offices/ocie/risk-alert-bd-controls-structured-securities-products.pdf> and FINRA disciplinary actions available at <http://www.finra.org//industry/disciplinary-actions/finra-disciplinary-actions-online?search=%20reverse%20convertibles>.

Mullainathan, Shafir, and Zinman, 2010; Gurun, Matvos, and Seru, 2016; Hastings, Hortaçsu, and Syverson, 2017). In the context of YEPs, the product would appear more attractive if an investor overweights the product coupon rate in her decision, and underweights the size or probability of the possible loss.

Consistent with such framing effects, I find the coupon rate is saliently displayed in product prospectuses, often even in the name of the product. Information about the payment at maturity when investors may lose some or all of their principal is less salient. The information is typically displayed only after the coupon rate and does not come with any indication of the size or the probability of losses, because these variables need to be estimated using option pricing techniques.

More interestingly, the possible loss of principal at maturity is often not emphasized as a loss but instead reframed as the physical delivery of the underlying. For example, the prospectus of the example product phrases the losses as "delivery of the reference stock instead of the principal amount," where the amount of the stock equals the face value of the product divided by the initial underlying price.²² This phrasing allows brokers to frame the payoffs of YEPs in terms of two mental accounts—holding the stock and getting additional coupon income. At least since Black (1975) and Shefrin and Statman (1993), brokers have been recognized using such framing manipulations, which give a misguided impression of abnormal gains. I screen payoff descriptions of all U.S. structured products and find physical delivery is almost exclusively used only in the loss domain and rarely in the gain domain. This pattern is consistent with banks using physical delivery of the underlying for framing and for making YEPs look more attractive.

Another mechanism to make the products look appealing is through underlying selection. The higher the volatility of the underlying, the higher the probability that investors will lose some of their invested capital. Therefore, strategic selection of highly volatile stocks can lead an investor to underestimate the probability of a loss if she underestimates the volatility. I find that the stocks used as underlyings have significantly higher implied volatility than possible alternative stocks. In a sample of the 750 largest

²²The prospectus is available at <https://www.rbccm.com/assets/rbccm/docs/expertise/equities/imported/usstructurednotes/file-612566.pdf>.

non-financial U.S. stocks, those used as an underlying have an implied volatility of at-the-money call options of 39.8% compared to 32.4% volatility in the rest of the sample. Table D.5 in the online appendix presents results from probabilistic regressions of underlying selection and shows that stocks with high implied volatility are significantly more likely used as underlying even after controlling for market capitalization, past returns, and trading volume. The relationship holds even after including underlying fixed effects, consistent with the hypothesis that banks are more likely to use a stock as an underlying at times when its implied volatility is high.

In summary, the evidence in this sections is consistent with the hypothesis that unsophisticated investors are more likely to view the products as more valuable because of framing, salient presentation of attractive product features, and strategic selection of underlyings. This evidence is consistent with the explanation that investors in YEPs may be unaware of their high costs and negative returns.

7.3 Marketing of YEPs

YEPs are issued by banks and distributed through broker-dealers who are compensated by commissions disclosed in the prospectus. In a sample of 25,241 products for which the platform reports the commissions, the average commission is 1.8%. Therefore the broker commissions account for about 50% of the estimated margins, and in their absence, the expected returns of the products would have been positive.

More importantly, these commissions are significantly higher than commissions from purchases of option contracts that approximate the payoffs of YEPs, such as the synthetic YEPs described in the previous section. One of the cheapest ways to execute the synthetic YEP position is to use low-commission online brokers. I estimate that under conservative assumptions, the average commissions in the sample of 17,000 synthetic YEPs is less than 0.20%, or around 10% of the YEP commission, when executed with online brokers.²³ Brokers affiliated with issuers of YEPs may have little incentives to direct investors to inexpensive online brokers, but are still compensated from op-

²³To estimate the commissions, I assume the synthetic YEP position is \$25,000 and the commission per option trade is \$15 and \$0.65 per contract. With a probability of 30%, the put options embedded in the synthetic position end up in the money, and the investor pays additional commissions for assignment /execution or to buy/sell options to close the original positions.

tion commissions charged by their brokerage company. I estimate that for three of the largest issuers in the sample ($N = 8,740$), the commissions from investing in synthetic YEP through the affiliated broker is 0.91%, or 47% of the respective YEP commission.²⁴ Therefore, brokers acting based on their best interest and not the clients' best interest, have strong incentives to steer investors to more expensive and inferior YEPs.

8 Conclusion

Innovation appears to proceed only very slowly in improving the design of retail financial products, likely because of mistakes of unsophisticated households (Campbell, 2006). One particular type of inefficiency emanates from cross-subsidies between naive and sophisticated households, which inhibits welfare-improving innovations (Gabaix and Laibson, 2006). YEPs analyzed in this paper represent another example of investment mistake. But in contrast to other markets, cross-subsidies between investors cannot play a role in the market for dominated YEPs, because any investor would be better off investing in listed options. My estimated fees imply that YEPs are overpriced to such an extent that their expected and realized returns are negative. The high costs of the products are often hidden or hard to understand, their bad performance is masked by biased performance measures used in the industry, and their design appears to cater to investor biases and the conflicted interest of brokers. Taken together, the evidence in this paper suggests rent-seeking by financial intermediaries is what inhibits welfare-improving financial innovation.

Yet, one can argue that YEPs represent only a small fraction of the retail financial markets. Why should we care? In other words, are cases of inferior financial products closer to the rule rather than an exception? On one hand, YEPs clearly differ from other retail financial products. YEPs are typically not listed, nor are they covered by independent research firms such as Morningstar, both of which likely foster competition and transparency among other retail financial products. Perhaps we should not be concerned.

²⁴To calculate the commissions, I use commissions per option trade of \$25, \$28.95, and \$62–80.75 for JPMorgan Chase, RBC, and UBS, respectively. The respective commissions per contract are \$0.65, \$1.5, and \$7.5. UBS charges an additional commission of 1.75% - 2.25% of the option principal.

On the other hand, however, many factors support a less optimistic view. The products are suitable only to relatively rich investors, they are relatively easy to compare with listed options on which information is easily accessible, and they have existed for long enough to allow investors to learn about their poor performance. The fair value disclosure mandated by the SEC gave investors a chance to better understand the embedded costs, and the high coupons of the products should raise a red flag about their high risk to any moderately experienced investor.

References

- Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel, 2015. Regulating consumer financial products: Evidence from credit cards. *Quarterly Journal of Economics* 130, 111–164.
- Aguilar, L., 2015. Regulators working together to serve investors - Speech at the Annual NASAA/SEC 19(d) Conference.
- Allen, F., 2012. Trends in financial innovation and their welfare impact: An overview. *European Financial Management* 18, 493–514.
- Allen, F., and D. Gale, 1994. *Financial innovation and risk sharing*. MIT Press.
- An, B., A. Ang, T. Bali, and N. Cakici, 2014. The joint cross section of stocks and options. *Journal of Finance* 69, 2279–2337.
- Anagol, S., and H. Kim, 2012. The impact of shrouded fees: Evidence from a natural experiment in the Indian mutual funds market. *American Economic Review* 102, 576–593.
- Andreasen, J., and B. Huge, 2011. Volatility interpolation. *Risk Magazine* 24, 86–89.
- Ausubel, L., 1991. The failure of competition in the credit card market. *American Economic Review* 81, 50–81.
- Barracrough, K., and R. Whaley, 2012. Early exercise of put options on stocks. *Journal of Finance* 67, 1423–1456.
- Baule, R., 2011. The order flow of discount certificates and issuer pricing behavior. *Journal of Banking and Finance* 35, 3120–3133.
- Ben-David, I., J. Li, A. Rossi, and Y. Song, 2019. What do mutual fund investors really care about?, Unpublished working paper.
- Bergstresser, D., 2008. The retail market for structured notes: Issuance patterns and performance, 1995-2008, Unpublished working paper.
- Bertrand, M., and A. Morse, 2011. Information disclosure, cognitive biases, and payday borrowing. *Journal of Finance* 66, 1865–1893.
- Bertrand, M., D. Karlan, S. Mullainathan, E. Shafir, and J. Zinman, 2010. What’s advertising content worth? Evidence from a consumer credit marketing field experiment. *Quarterly Journal of Economics* 125, 263–306.
- Beshears, J., J. Choi, D. Laibson, and B. Madrian, 2011. How does simplified disclosure affect individuals’ mutual fund choices? In *Explorations in the economics of aging*, ed. D. Wise, Chapter 2, 75–96. Chicago:University of Chicago Press.

- Bessembinder, H., 2018. Do stocks outperform treasury bills? *Journal of Financial Economics* 129, 440–457.
- Black, F., 1972. Capital market equilibrium with restricted borrowing. *Journal of Business* 45, 444–455.
- Black, F., 1975. Fact and fantasy in the use of options. *Financial Analysts Journal* 31, 36–41.
- Bordalo, P., N. Gennaioli, and A. Shleifer, 2016. Competition for attention. *Review of Economic Studies* 83, 481–513.
- Broadie, M., M. Chernov, and M. Johannes, 2009. Understanding index option returns. *Review of Financial Studies* 22, 4493–4529.
- Burth, S., T. Kraus, and H. Wohlwend, 2003. The pricing of structured products in the Swiss market. *Journal of Derivatives* 11, 55–69.
- Calvet, L., C. Célérier, P. Sodini, and B. Vallée, 2017. Can financial innovation solve household reluctance to take risk?, Unpublished working paper.
- Campbell, J., 2006. Household finance. *Journal of Finance* 61, 1553–1604.
- Campbell, J., 2016. Restoring rational choice: The challenge of consumer financial regulation. *American Economic Review: Papers and Proceedings* 106, 1–30.
- Campbell, J., H. Jackson, B. Madrian, and P. Tufano, 2011. Consumer financial protection. *Journal of Economic Perspectives* 25, 91–114.
- Carlin, B., 2009. Strategic price complexity in retail financial markets. *Journal of Financial Economics* 91, 278–287.
- Carlin, B., S. Kogan, and R. Lowery, 2013. Trading complex assets. *Journal of Finance* 68, 1937–1960.
- Célérier, C., and B. Vallée, 2017. Catering to investors through security design: Headline rate and complexity. *Quarterly Journal of Economics* 132, 1469–1508.
- Coval, J., J. Jurek, and E. Stafford, 2009. The economics of structured finance. *Journal of Economic Perspectives* 23, 3–25.
- Deng, G., T. Dulaney, T. Husson, C. McCann, and M. Yan, 2015. Ex-post structured product returns: Index methodology and analysis. *Journal of Investing* 24, 45–58.
- Di Maggio, M., and M. Kacperczyk, 2017. The unintended consequences of the zero lower bound policy. *Journal of Financial Economics* 123, 59–80.
- Duarte, F., and J. Hastings, 2012. Fettered consumers and sophisticated firms: Evidence from Mexico’s privatized social security market, NBER Working Paper No. 18582.

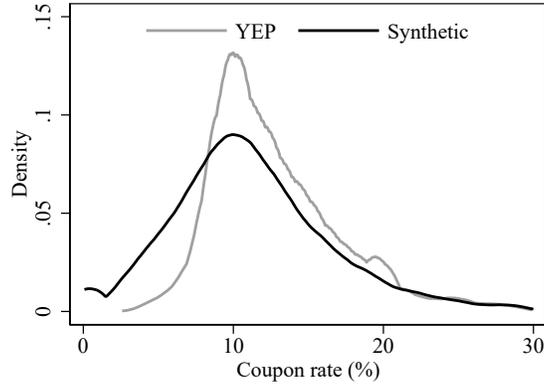
- Dupire, B., 1994. Pricing with a smile. *RISK* 7, 18–20.
- Egan, M., 2019. Brokers versus retail investors: Conflicting interests and dominated products. *Journal of Finance* 74, 1217–1260.
- Elul, R., 1995. Welfare effects of financial innovation in incomplete markets economies with several consumption goods. *Journal of Economic Theory* 65, 43–78.
- Fama, E., and K. French, 2015. A five-factor asset pricing model. *Journal of Financial Economics* 116, 1–22.
- Frame, W., L. Wall, and L. White, 2018. Technological change and financial innovation in banking: Some implications for Fintech, FRB Atlanta Working Paper.
- Gabaix, X., and D. Laibson, 2006. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics* 121, 505–541.
- Gennaioli, N., A. Shleifer, and R. Vishny, 2012. Neglected risks, financial innovation, and financial fragility. *Journal of Financial Economics* 104, 452–468.
- Gennaioli, N., A. Shleifer, and R. Vishny, 2014. Money doctors. *The Journal of Finance* 70, 91–114.
- Ghent, A., W. Torous, and R. Valkanov, 2019. Complexity in structured finance. *Review of Economic Studies* 86, 694–722.
- Goetzmann, W., J. Ingersoll, M. Spiegel, and I. Welch, 2007. Portfolio performance manipulation and manipulation-proof performance measures. *Review of Financial Studies* 20, 1503–1546.
- Greenwood, R., and D. Scharfstein, 2013. The growth of finance. *Journal of Economic Perspectives* 27, 3–28.
- Gurun, U., G. Matvos, and A. Seru, 2016. Advertising expensive mortgages. *Journal of Finance* 71, 2371–2416.
- Hanson, S., and J. Stein, 2015. Monetary policy and long-term real rates. *Journal of Financial Economics* 115, 429–448.
- Hastings, J., A. Hortaçsu, and C. Syverson, 2017. Sales force and competition in financial product markets: The case of Mexico’s social security privatization. *Econometrica* 85, 1726–1761.
- Heidhues, P., B. Kőszegi, and T. Murooka, 2016. Exploitative innovation. *American Economic Journal: Microeconomics* 8, 1–23.

- Henderson, B., and N. Pearson, 2011. The dark side of financial innovation: A case study of the pricing of a retail financial product. *Journal of Financial Economics* 100, 227–247.
- Henderson, B., N. Pearson, and L. Wang, 2020a. Retail derivatives and sentiment: A sentiment measure constructed from issuances of retail structured equity products, Unpublished working paper.
- Henderson, B., N. Pearson, and Li Wang, 2020b. Pre-trade hedging: Evidence from the issuance of retail structured products. *Journal of Financial Economics* 137, 108–128.
- Hortaçsu, A., and C. Syverson, 2004. Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds. *Quarterly Journal of Economics* 119, 403–456.
- Hull, J., and A. White, 2013. LIBOR vs. OIS: The derivatives discounting dilemma. *Journal of Investment Management* 11, 14–27.
- Koijen, R., and M. Yogo, 2017. The fragility of market risk insurance, Unpublished working paper.
- Lerner, J., and P. Tufano, 2011. The consequences of financial innovation: A counterfactual research agenda. *Annual Review of Financial Economics* 3, 41–85.
- Levenshtein, V., 1966. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics-Doklady* 10, 707–710.
- Li, X., A. Subrahmanyam, and X. Yang, 2018. Can financial innovation succeed by catering to behavioral preferences? Evidence from a callable options market. *Journal of Financial Economics* 128, 38–65.
- Lian, C., Y. Ma, and C. Wang, 2019. Low interest rates and risk taking: Evidence from individual investment decisions. *Review of Financial Studies* 32, 2107–2148.
- Martin, I., 2017. What is the expected return on the market? *Quarterly Journal of Economics* 132, 367–433.
- Philippon, T., 2015. Has the US finance industry become less efficient? On the theory and measurement of financial intermediation. *American Economic Review* 105, 1408–1438.
- Philippon, T., 2016. The FinTech opportunity, NBER Working paper.
- Ross, S., 1976. Options and efficiency. *Quarterly Journal of Economics* 1, 75–89.
- Schoar, A., 2012. Comment on Lerner, J. and P. Tufano, The consequences of financial innovation: A counterfactual research agenda. In *The rate and direction of inventive activity revisited*, ed. J. Lerner and S. Stern, Chapter 11, 576–578. Chicago:University of Chicago Press.

- Shefrin, H., and M. Statman, 1993. Behavioral aspects of the design and marketing of financial products. *Financial Management* 22, 123–134.
- Stein, J., 2013. Overheating in credit markets: origins, measurement, and policy responses, Federal reserve board of governors working papers, Federal Reserve Board of Governors.
- Szymanowska, M., J. Horst, and C. Veld, 2009. Reverse convertible bonds analyzed. *Journal of Futures Markets* 29, 895–919.
- Tufano, P., 2003. Financial innovation. In *Handbook of the economics of finance*, ed. G. Constantinides, M. Harris and R. Stulz, 307–336.
- Volcker, P., 2009. Think more boldly. *Wall Street Journal*.
- Zingales, L., 2015. Presidential Address: Does finance benefit society? *Journal of Finance* 70, 1327–1363.

Figures and Tables

Panel A: Coupon rates of YEPs and synthetic YEPs



Panel B: Excess coupon rate of synthetic YEPs

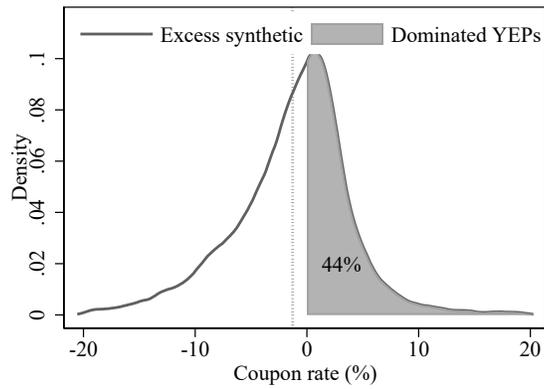


Fig. 2: Synthetic YEPs and dominated products

Panel A plots the distribution of annual coupon rates for YEPs and for their synthetic counterparts. Synthetic securities are constructed using put options and lending as described in Section 4. Panel B plots excess synthetic coupon rates calculated as the difference between the synthetic rate and YEP coupon rate. Products highlighted in gray represent products that are dominated by their synthetic counterparts, that is, products with positive excess synthetic rate. The vertical line in Panel B marks the mean excess rate.

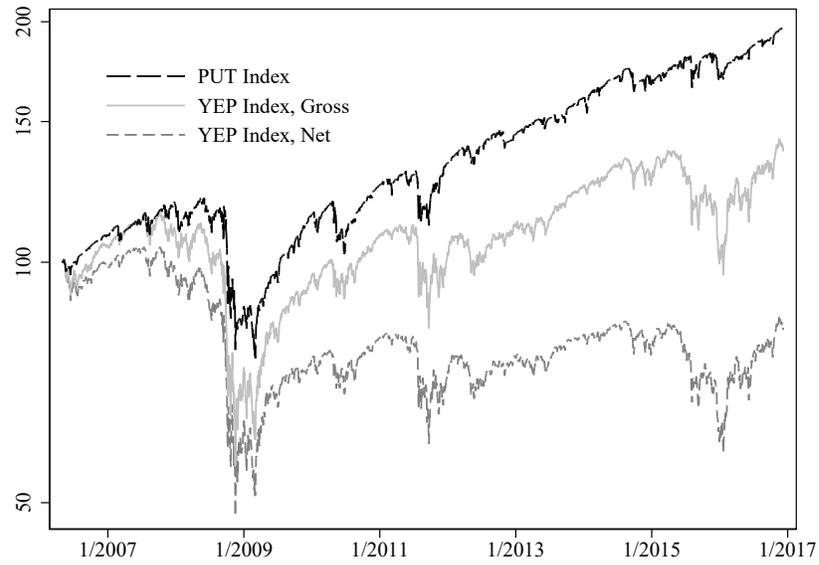
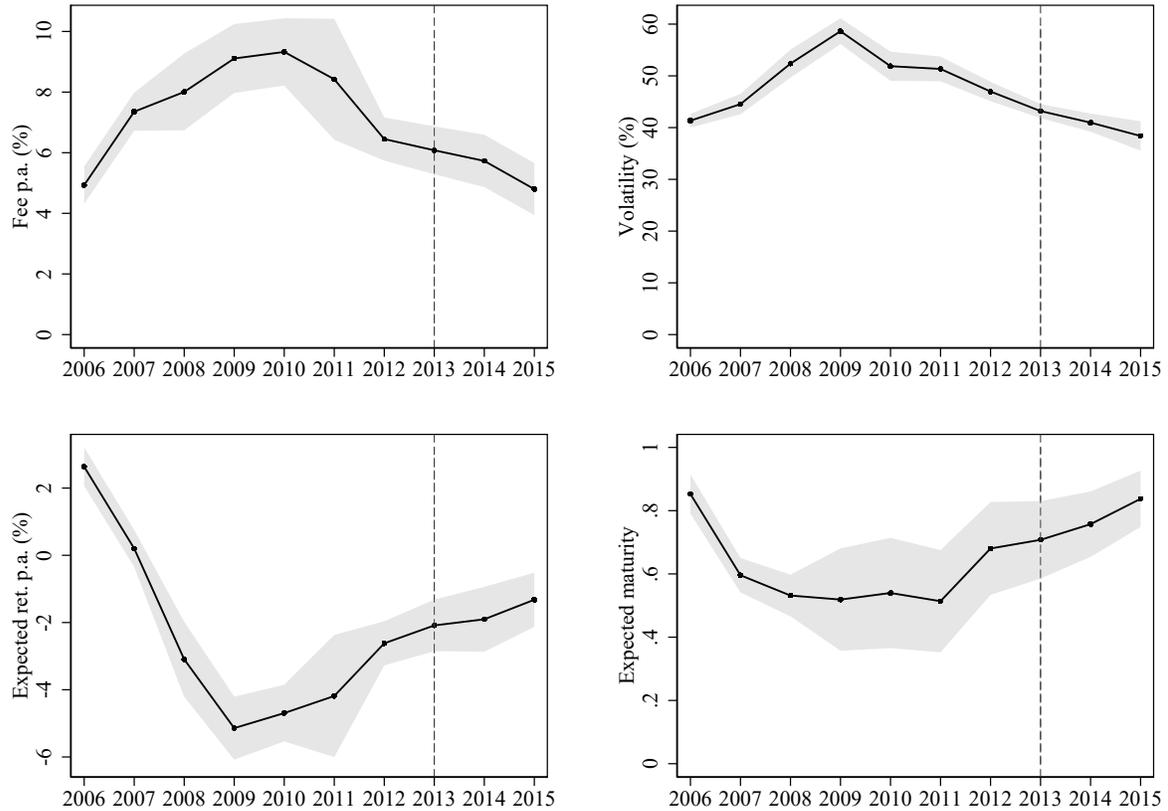


Fig. 3: YEP indexes

The figure plots the log index growth between May 2006 and December 2016 for YEP indexes with and without fees, and for the PUT index. The construction of the YEP indexes is described in Section 5.1. PUT Index is Cboe S&P 500 PutWrite Index.

Panel A: Full time series



Panel B: Fees around mandatory disclosure

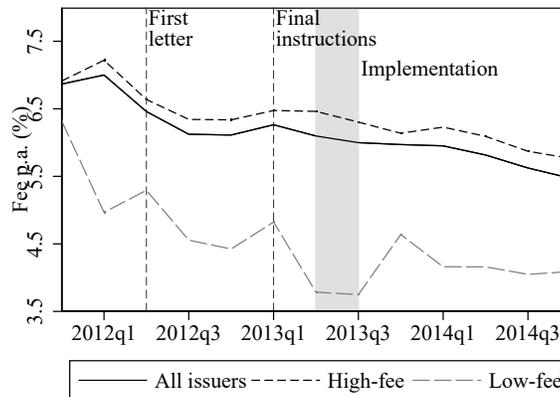


Fig. 4: Fees, expected returns, and product characteristics over time

The figure shows the annual averages and 95% confidence intervals (with standard errors clustered at the issuer level) of fees, expected returns, expected maturity, and volatility in Panel A, and quarterly averages of fees in Panel B. Expected returns are estimated assuming the expected excess return on the underlying equals $\mu - r = \beta \times 6\% p.a.$ The vertical lines in Panel A depicts the year when the SEC mandated disclosure of the issuer estimated fair value. Vertical lines in Panel B denote the quarter when the SEC sent the first letter about the mandatory disclosure, the quarter when final instructions for disclosure were published, and two quarters during which issuers implemented the disclosure.

Table 1: YEP example

The table presents an example of a product, its payoff translation, and synthetic replication. Prices of the underlying in the payoff description and translation are normalized to 100 at issuance.

Panel A: Product characteristics

Name	Reverse Convertible Notes No. 3296
CUSIP	78008TYE8
Volume	\$0.437 million
Initial strike date	Nov 23, 2011
Term	3 months
Coupon rate p.a.	11.5%
Underlying	JPMorgan Chase & Co.
Commission	1.75%

Panel B: Payoff translation

Payoff description This is an income product linked to the share of JPMorgan Chase. The product offers a coupon of 11.50% p.a., paid monthly throughout the investment period. At maturity, the product offers a capital return of 100%, if the final share level is equal to or greater than the initial level or if the underlying does not fall by more than 25% from its initial level at any time during the investment. If the underlying does fall by more than 25% from its initial level at any time during the investment and the final share level is lower than the initial level, the capital return equals 100%, minus 1% for every 1% fall, paid in cash or in shares.

Translated formula

$$P_T = \begin{cases} 100 & \text{if } S_t \geq 75 \quad \forall t = 1, \dots, T \\ S_T & \text{otherwise} \end{cases}$$

$$P_m = 11.5 \times \Delta t \quad \text{for monthly observation dates } m = 1, 2, 3$$

Panel C: Synthetic replication

Synthetic coupon	26.7%
Deposit income	0.1%
Option income	$26.6\% = (n_1 p_1 - n_2 p_2) / (S_0 D)$,
	where option income is generated by short and long positions in put options:
	Short leg:
	Strike price: $K_1 = \$21.00$
	Quantity: $n_1 = 8.38$
	Price: $p_1 = \$0.82$
	Long leg
	Strike price: $K_2 = \$20.00$
	Quantity: $n_2 = 7.38$
	Price: $p_2 = \$0.69$
	Day count using the expiry date of the options, Feb 18, 2012, is $D = 85/360$.
	Face value equals the initial price of the underlying, $S_0 = \$28.38$.

Table 2: Summary statistics

The table reports summary statistics of product characteristics at issuance (Panel A), and inputs to the pricing model (Panel B). *Volume* is issuance volume in million \$. *Coupon rate* is the product annual coupon rate extracted from its payoff description. If the price of the underlying falls below the *Barrier* level (in % of the initial underlying price), investors could lose some or all of the principal value. *Maximum term* (in years) is the maximum maturity of a product if it does not terminate early. Δ is the estimated product delta at issuance. Broker's *Commission* is in % of product price. The sample consists of 28,383 YEPs issued between January 2006 and September 2015.

Panel A: Product characteristics

	Mean	Vol.-wtd. average	Std. Dev.	p1	p25	p75	p99	Observations
Volume	2.0	–	5.3	0.0	0.2	2.0	20.2	28,383
Coupon rate	12.8	11.8	4.7	5.8	9.6	15.0	28.7	28,383
Barrier	73.6	75.5	7.6	50.0	70.0	80.0	90.0	28,383
Maximum term	0.8	1.0	0.7	0.2	0.5	1.0	5.0	28,383
Commission	1.8	1.8	0.8	0.5	1.5	2.0	5.0	25,241

Panel B: Underlying factor loadings

	Mean	Vol.-wtd. average	Std. Dev.	p1	p25	p75	p99	Observations
$\hat{\beta}$	1.6	1.6	0.8	0.3	1.1	2.0	3.9	27,289
$\hat{\beta}_{SMB}$	0.4	0.3	1.0	-2.0	-0.2	0.9	3.6	27,289
$\hat{\beta}_{HML}$	0.2	0.3	1.4	-3.9	-0.7	0.9	3.9	27,289
$\hat{\beta}_{CMA}$	-0.9	-0.8	1.8	-5.8	-1.9	0.2	4.2	27,289
$\hat{\beta}_{RMW}$	-0.2	-0.1	1.6	-5.6	-1.0	0.6	3.7	27,289

Table 3: Margins and embedded fees

The table reports estimates of product margins (Panel A) and annual fees (Panel B) obtained from the pricing model described in Section 3.1. The first two columns of Panel A present the values not adjusted for credit risk for the whole sample (first column) and for the subsample with available CDS spreads (second column). The third column presents values adjusted for credit risk. *Margin* (in % of issue price) is the difference between the fair product value and the issue price. *Fees p.a.* (in %) are defined as the margin divided by the expected term of the product, where the expected term is calculated using risk-neutral probabilities of early termination. Standard errors are clustered at the issuer level and reported in brackets.

Panel A: Margins

	Adjusted for credit risk		
	No	No	Yes
	Full Sample	CDS available	CDS available
Mean	3.63 (0.19)	3.76 (0.18)	3.90 (0.19)
Vol.-wtd. Average	3.53	3.57	3.78
p25	2.48	2.69	2.80
p50	3.62	3.75	3.90
p75	4.62	4.71	4.88
Observations	28,383	21,617	21,617

Panel B: Fees p.a.

	Full Sample	Expected term		
		2 – 4 months	4 – 8 months	> 8 months
Mean	7.18 (0.51)	11.73 (0.33)	7.18 (0.32)	3.97 (0.16)
Vol.-wtd. Average	5.73	10.69	6.43	3.81
p25	3.82	6.69	5.07	2.79
p50	6.13	11.11	7.07	3.95
p75	9.13	16.65	8.91	5.04
Observations	28,383	6,389	12,932	9,062

Table 4: Expected returns

The table reports estimates of net-of-fees expected returns of the products. The column labels indicate the model used to estimate the expected return on the underlying. $\hat{\beta}$ is the CAPM beta. \overline{CRSP}_t is the value-weighted average. $SVIX$ is the 1-year equity premium based on the SVIX index (Martin, 2017). $\hat{\beta}$ is a vector of Fama and French (2015) factor loadings, and $\overline{FF5}_t$ is a vector of the respective mean factor values. Betas are estimated using 24 – 60 monthly returns preceding the initial valuation date. Average factor returns are over the period from January 1996 until the last month before the initial valuation date of the product. Panel A reports summary statistics of underlying factor loadings, Panel B reports total returns, and Panel C reports annualized expected returns. The sample consists of 27,289 products issued between January 2006 and September 2015. Standard errors are clustered at the issuer level and reported in brackets.

Panel A: Total expected returns

	Expected excess underlying return				
	$\hat{\beta} \times 6\% \text{ p.a.}$	$\hat{\beta} \times 8\% \text{ p.a.}$	$\hat{\beta} \times \overline{CRSP}_t$	$\hat{\beta} \times SVIX_t$	$\hat{\beta} \cdot \overline{FF5}_t$
Mean	-0.87 (0.14)	-0.29 (0.14)	-0.41 (0.14)	-1.10 (0.23)	-1.04 (0.20)
Vol.-wtd. average	-0.18	0.47	0.34	-0.65	-0.20
p25	-2.27	-1.78	-1.89	-2.58	-3.17
p50	-1.1	-0.54	-0.66	-1.13	-1.09
p75	0.34	0.96	0.83	0.37	0.89
Observations	27,289	27,289	27,289	15,759	27,289

Panel B: Expected returns p.a.

	Expected excess underlying return				
	$\hat{\beta} \times 6\% \text{ p.a.}$	$\hat{\beta} \times 8\% \text{ p.a.}$	$\hat{\beta} \times \overline{CRSP}_t$	$\hat{\beta} \times SVIX_t$	$\hat{\beta} \cdot \overline{FF5}_t$
Mean	-2.61 (0.33)	-1.63 (0.32)	-1.84 (0.33)	-3.29 (0.44)	-2.86 (0.34)
Vol.-wtd. average	-1.09	-0.11	-0.30	-1.79	-1.12
p25	-5.03	-4.06	-4.29	-6.33	-6.58
p50	-2.03	-1.03	-1.24	-2.37	-2.19
p75	0.54	1.57	1.37	0.65	1.58
Observations	27,289	27,289	27,289	15,759	27,289

Table 5: Ex-post returns

The table reports estimates of product ex-post returns. *Product return* is the sum of the payoff at maturity and all coupon payments. *Benchmark return* is the cumulative return of delta equivalent daily adjusted positions in the underlying equity and risk-free rate at maturity. *Abnormal return* is the difference between product return and benchmark return. The sample consists of 27,578 products issued between January 2006 and September 2015 and maturing before January 2018. Standard errors are clustered at the issuer level and reported in brackets.

Product returns

	Total returns			Returns p.a.		
	Product	Benchmark	Abnormal	Product	Benchmark	Abnormal
Mean	-3.80 (1.27)	-0.45 (0.73)	-3.36 (0.63)	-2.62 (2.46)	4.74 (1.80)	-7.36 (0.73)
Vol.-wtd. average	-3.98	-0.74	-3.25	-3.46	2.47	-5.93
p25	-4.67	-0.67	-5.63	-7.01	-1.04	-11.60
p50	3.90	5.24	-2.39	10.10	11.65	-5.20
p75	7.00	8.95	-0.18	13.68	20.53	-0.34
Observations	27,578	27,578	27,578	27,578	27,578	27,578

Table 6: Dominated products

The table reports the fraction of products that are dominated by their synthetic counterpart constructed using put options and lending as described in Section 4. The first column uses the closest option expiration as the maturity of synthetic YEP and includes income from lending. The second column excludes lending income. The third column uses the exact YEP maturity as the maturity of synthetic YEP. The sample in columns 1 and 2 consists of 17,000 products issued between January 2006 and September 2015 that can be statically approximated with up to two positions in put options and lending. In column 3, the sample is restricted to YEPs with existing listed options that expire up to 40 days before the maturity of YEPs.

Fraction of dominated products

Specification	Baseline	Without lending income	Exactly matching maturity
Equal weighted	43.60	33.88	27.90
Vol.-weighted	44.70	30.75	30.37
Observations	17,000	17,000	8,915

Table 7: YEP indexes

Panel A reports summary statistics of daily excess returns of YEP indexes described in Section 5.1. Mean and standard deviations are annualized. Newey-West standard errors with 20 lags are in parentheses. Panel B reports four-factor regressions of daily excess returns of YEP indexes. Alphas are annualized. Newey-West t -statistics with 20 lags are in parentheses.

Panel A: Index summary statistics

	Mean	ex-	S.E.	S.D.	Daily skew	Daily	kur-	Observations
	cess	cess				tosis		
	return	return						
PUT	6.36		(3.50)	14.25	-0.19	30.06		2,688
YEP Gross	3.93		(5.29)	19.29	-0.18	18.18		2,688
YEP Net	-0.90		(5.30)	19.24	-0.18	18.19		2,688

Panel B: Index factor loadings

Dep. var.	YEP Gross	YEP Gross	YEP Net	YEP Net
PUT	1.15	1.05	1.15	1.04
	(29.27)	(38.85)	(29.38)	(38.64)
SMB		0.24		0.24
		(5.45)		(5.44)
HML		-0.01		-0.01
		(-0.26)		(-0.22)
Momentum		-0.31		-0.31
		(-7.07)		(-7.02)
Constant	-3.38	-3.12	-8.19	-7.91
	(-1.15)	(-1.24)	(-2.75)	(-3.11)
Observations	2,688	2,688	2,688	2,688
R^2	0.72	0.80	0.72	0.80

Table 8: Fair value disclosure

Panel A reports summary statistics of the fair product values estimated using the pricing model described in Section 3.1 (*model value*) and the estimated values disclosed by the issuers (*issuer value*) following the 2013 introduction of fair value disclosure. Both values are reported as a fraction of the issue price in %. The sample consists of 6,372 YEPs for which both the model value and the issuer value are available. Panel B reports coefficient estimates and *t*-statistics from versions of the regression:

$$y_{ikt} = \beta \text{High-Fee}_k \times \text{Post}_t + \text{Post}_t + \lambda_k + \epsilon_{ikt},$$

where y_{ikt} is either annualized fee, margin, or volume sold of product i issued in quarter t by issuer k , High-Fee_k is an indicator equal to 1 for issuers with above-median fees in year 2012, Post_t is an indicator equal to 1 at any point after 2013q2, and λ_k stands for issuer fixed effects. The sample covers 7,558 products issued in 2012 or between 2013q3 and 2014q2. *t*-statistics are based on standard errors clustered at the issuer level.

Panel A: Summary statistics of estimated values

	Mean	Vol.-wtd. average	S. D.	p1	p25	p75	p99	Observations
Model value	96.42	96.78	1.27	93.41	95.63	97.21	99.75	6,372
Issuer value	96.61	96.92	0.91	94.14	96.10	97.20	98.61	6,372

Panel B: Effect of fair value disclosure

Dep. Var.	Fee p.a.	Fee p.a.	Margin	Margin	Volume	Volume
Post \times High-Fee		0.63 (0.88)		-0.06 (-0.38)		-0.97 (-1.10)
Post	-0.36 (-1.41)	-0.91 (-1.40)	-0.07 (-1.14)	-0.01 (-0.09)	0.38 (2.23)	1.23 (1.41)
Issuer FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7,558	7,558	7,558	7,558	7,558	7,558

Online appendix for:
Engineering lemons

Petra Vokata[†]
Ohio State University

October 2020

[†]Ohio State University, 832 Fisher Hall 2100 Neil Avenue Columbus, OH 43210. E-mail: vokata.1@osu.edu.

A Data Appendix

A.1 Platform Coverage

To confirm database is comprehensive, I validate its coverage with the filings in the EDGAR system. I download all prospectuses filed as 424B2 forms between January 2002 and September 2015 and select all files that contain a CUSIP code and the keyword "linked," and that do not include the keyword "ETN." I search the prospectuses for the list of CUSIP codes recorded by the platform. Table A.1 shows the results of this exercise and indicates the platform has lower coverage before the year 2006 and excellent coverage thereafter. Closer examination reveals the prospectuses that do not match with any CUSIP code do not cover structured products but mostly fixed-rate notes. In addition, the total issuance volume reported by the platform for products issued in 2014 aligns with the aggregate volume of structured notes reported by Bloomberg. Therefore, I believe the platform coverage of structured products from 2006 onward is exhaustive.

Table A.1: Platform Coverage

The table reports the estimated platform coverage. I start with all prospectuses filed as 424B2 forms in the EDGAR system between January 2002 and September 2015 and select all prospectuses that contain a CUSIP code and the keyword "linked" and that do not include keyword "ETN." *All prospectuses* reports the number of files that fit these criteria each year. Next, I search these prospectuses for a list of all CUSIP codes recorded by the platform. *Matched with platform* reports the number of these matched files. *Estimated coverage* is the number of matched files divided by the total number of files.

Year	All prospectuses	Matched with platform	Estimated coverage (%)
2002	46	3	7
2003	95	1	1
2004	131	-	
2005	136	13	10
2006	1,201	840	70
2007	3,893	2,840	73
2008	4,500	3,699	82
2009	3,176	2,923	92
2010	5,073	4,440	88
2011	7,571	6,959	92
2012	9,250	8,941	97
2013	9,729	9,386	96
2014	10,804	10,487	97
2015	8,854	8,266	93
2006–15	64,051	58,781	92

Table A.2: Final Sample by Year

The table reports the number of products and total issuance volume in the final translated sample by issuance year. The sample consists of 28,383 YEPs issued between January 2006 and September 2015.

Issuance Year	Number of products	Issuance volume (million \$)
2006	813	2,404
2007	3,506	8,502
2008	3,339	7,263
2009	1,995	4,292
2010	3,045	6,778
2011	3,135	6,853
2012	3,955	4,961
2013	3,354	5,110
2014	3,576	6,336
2015	1,665	3,176

A.2 Sample Selection

The platform covers 36,742 structured products (\$124.841 billion) issued in the U.S. as registered notes between January 2006 and September 2015 (time of data retrieval) and categorized as "income" products. The list below describes the criteria applied to construct the dataset of YEPs and the number in product and issuance volume dropped due to each criterion.

- (a) (8 products, \$0.119 bn) I exclude products categorized by the platform as "Private Banking" products.
- (b) (4,994 products, \$61.039 bn) I drop products linked to non-equity asset classes for which data on returns and implied volatilities are not available.
- (c) (1,103 products, \$2.373 bn) I exclude products with incomplete data for valuation: products that cannot be reliably matched to OptionMetrics IvyDB US or products with incomplete payoff information.

I use the resulting sample of 30,637 products (\$61.309 billion) to train the translation algorithm. The translated sample covers the most frequent payoffs, 28,383 products, and \$55.675 billion of issuance volume. Table A.2 lists the number of products and issuance volume by issuance year and Tables A.3 and A.4 present the largest issuers and the most frequent underlyings covered by the sample.

A.3 Payoff Validation

To validate the accuracy of the translated payoff formulas, I use three independent data sources. First, I validate the calculated ex-post returns with the ex-post returns reported by the platform.

Table A.3: Issuers

The table reports the 15 largest issuers in the sample, the number of products they issued, and issuance volume. The sample consists of 28,383 YEPs issued between January 2006 and September 2015.

Issuer	Number of products	Issuance volume (million \$)
Barclays Bank	5,076	13,430
UBS	8,879	7,670
JPMorgan Chase	3,377	6,122
RBC	4,447	5,298
Citigroup Funding	282	5,199
Morgan Stanley	492	3,871
ABN Amro Bank	1,754	2,276
HSBC Bank	629	2,005
Deutsche Bank	501	1,869
Bank of America	56	1,770
SG Structured Products	476	1,369
Eksportfinans ASA	381	1,182
Svensk Exportkredit	24	613
Credit Suisse	302	515
Rabobank	188	459

The advantage of this validation is that the ex-post returns can be calculated with high precision and one can therefore detect even relatively small deviations. The disadvantage of this approach is that the platform covers only a subsample of ex-post returns and has a relatively small coverage for autocallable products. More importantly, because platform analysts calculate ex-post returns from the product characteristics recorded in the platform, comparison with platform returns will not detect errors in the initial entry of product terms.

To address these concerns, I also compare the calculated ex-post returns with an independent consulting firm that publishes approximate product returns on its website.¹ The advantage of this comparison is again the fact that the returns can be calculated with high precision. The disadvantage is again low coverage of autocallable products even by the consulting firm.

Finally, I also compare my estimates of product fair values post issuance with the secondary prices reported in TRACE. Although the reported TRACE prices include markdowns and therefore the comparison is less precise, they cover all types of product payoffs in my sample and allow me to validate translated formulas even for autocallable products. In addition, because TRACE prices are recorded directly by brokers, they are less prone to possible errors in calculating returns that may be present both in the platform and consulting firm returns.

¹I extracted the returns from product reports available at <https://www.slcg.com/pdf/tearsheets/>.

Table A.4: Underlying Equities

The table reports the 40 most frequent underlying assets in the sample, the number of products that are linked to them, and their total issuance volume.

Underlying	Number of products	Issuance volume (million \$)
Apple	1,251	3,685
United States Steel	722	829
Freeport-McMoRan	711	1,234
Bank of America	458	1,170
Facebook	439	689
Peabody Energy	389	575
Ford	386	1,722
JPMorgan Chase	363	1,331
Las Vegas Sands	335	629
General Electric	329	1,401
Caterpillar	324	581
Chesapeake Energy	286	655
Halliburton	268	793
General Motors	252	801
Valero Energy	250	708
Alcoa	245	586
MetLife	241	726
Micron Technology	240	216
Genworth Financial	239	266
Amazon	238	466
Citigroup	228	367
Yahoo	223	431
BlackBerry	223	307
Netflix	222	261
United Rentals	215	219
First Solar	208	199
PulteGroup	205	91
Alpha Natural Resources	204	145
Wells Fargo	200	703
Deere & Company	200	566
Delta Air Lines	198	359
Silver Wheaton	198	235
Joy Global	196	231
Dow Chemical	195	517
Schlumberger	193	667
Cummins	190	310
Arch Coal	174	308
Morgan Stanley	174	228
Tesoro	170	224
SanDisk	163	298

I start my analysis with the comparison with TRACE prices because it reveals the most important issue with the payoff descriptions. In particular, I find 1,428 product-day-price observations with significant deviations between mine and TRACE prices among products with "Airbag" in the name. By comparing product terms in the prospectuses with the payoff descriptions recorded in the platform, I find the platform has an imprecise description of the product payoff at maturity in cases when the underlying drops below the barrier. That is, the original textual description includes errors in the level of the payoff drop at the barrier and in the payoff slope below the barrier. I correct these errors for all 1,280 "Airbag" products in my final sample. After the correction, TRACE prices are highly correlated with my estimates (99.1% correlation in a sample of 24,746 product-day-price observations). Figure A.1 plots my estimates against TRACE prices. For more details about mean differences, see Section C.1.

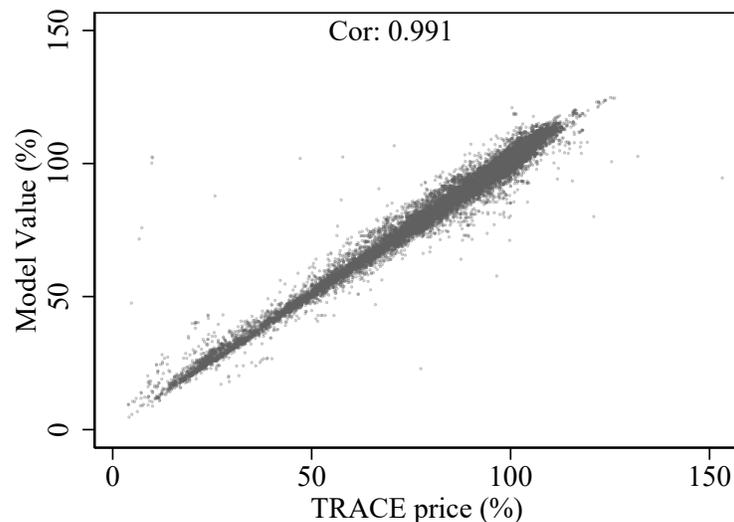


Figure A.1: Model values vs. secondary TRACE prices

The figure plots product fair values estimated using the pricing model described in Section 3.1 against secondary prices (in %) reported in TRACE. The sample consists of 24,746 product-day-price observations.

I next turn to validation with the platform returns. Because TRACE prices include variable markdowns and my fair values have to be estimated, they do not allow inspection of small discrepancies. Ex-post returns, on the other hand, can be calculated precisely. Figure A.2 plots the returns calculated from the translated formulas against total returns reported in the platform² as well as against annualized returns published by the consulting firm.

²I exclude "Airbag" products from this analysis due to their imprecise descriptions in the platform.

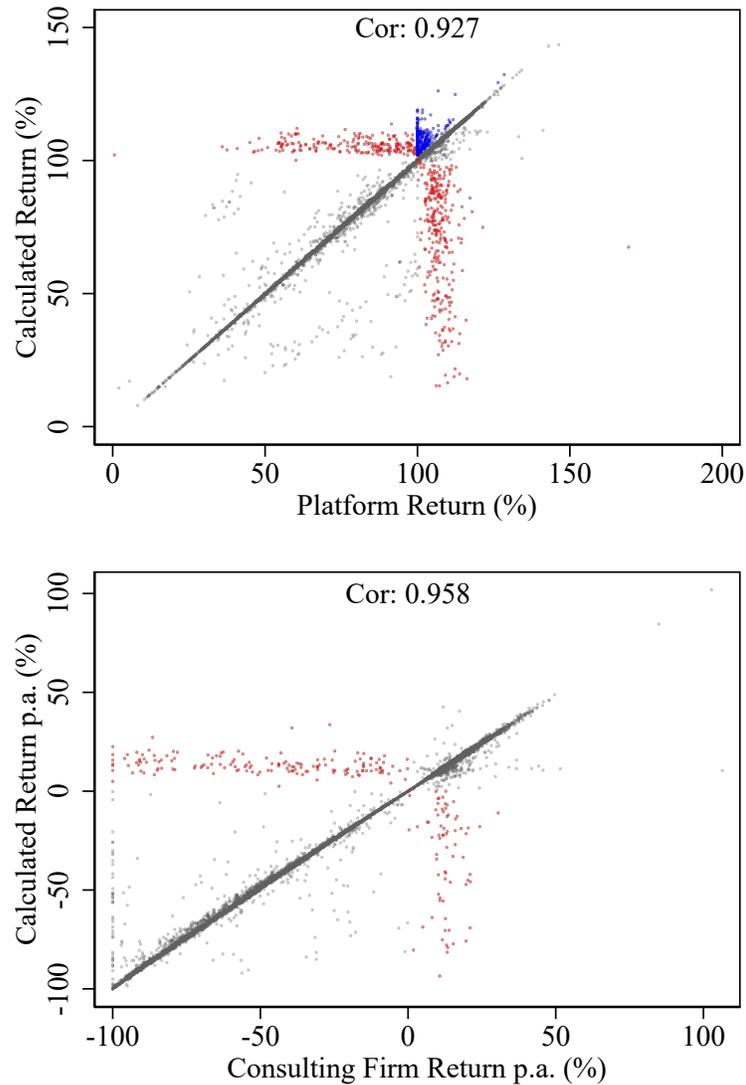


Figure A.2: Calculated Returns vs. Returns Reported by Platform and Consulting Firm

The figures plots returns calculated from translated payoff formulas against total returns reported by the platform and against annualized returns published by a consulting firm. The sample covers 14,439 products covered both by my sample and by the platform and 11,272 products covered by my sample and by the consulting firm.

Before discussing discrepancies between the returns, two features of the figures deserve comment. First, both figures display a high degree of correlation ($\rho = 0.927$ for platform returns and $\rho = 0.958$ for consulting firm returns) with the majority of the observations lying close to the 45-degree line. These patterns imply that the translated formulas are highly accurate and that the platform returns are arguably a good data source to analyze ex-post performance of the market. Second, comparison with the platform shows that, at least in the vast majority

of the cases, the platform returns include coupons paid over the life of the products, as do my returns calculated from the translated payoff formulas. C  lerier and Vall  e (2017) exclude from their analysis of ex-post performance of the European market products that pay coupons during the life of a product, because their data do not include coupon payment realization. I find the U.S. database of the platform includes coupon payments in product returns.

I now turn to the inspection of two types of return discrepancies that emerge from the figures. First, both figures show a number of cases (highlighted in red) in which the returns disagree about whether the product paid back the full principal at maturity ($N = 509$ for platform returns and $N = 228$ for consulting firm returns). Second, comparison with the platform returns shows a number of cases in which the platform calculates lower coupon payments than the translated formulas (highlighted in blue, $N = 292$).

The fact that the comparison with TRACE prices did not display the same differences in expected principal repayment at maturity already suggests the discrepancies are more likely driven by errors in platform or consulting firm returns. Nonetheless, I manually inspect the observations in red (for the platform) for the three most common underlyings: Freeport-McMoRan, Washington Mutual, and Deere & Company. I am able to track the majority of the discrepancies in products linked to Freeport-McMoRan and Deere & Company to their stock splits on February 2, 2011, and December 4, 2007, respectively. That is, the return in the platform differs because it appears to be based on non-adjusted prices, whereas my calculation correctly uses adjusted prices. The discrepancies in products linked to Washington Mutual appear to be related to the reassignment of "WM" ticker to Waste Management after Washington Mutual filed for Chapter 11 bankruptcy and was delisted.

As an additional check, I cross-validate platform discrepancies highlighted in red with the returns of the consulting firm in the top panel of Figure A.3. The idea of this exercise is to closely inspect cases in which two data sources disagree with my returns. For the majority of the cases in which mine and platform returns disagree, the consulting firm returns align with my returns. I again manually inspect the observations for which even the consulting firm returns show discrepancies for the three most common underlyings (Deere & Company, Freeport-McMoRan, and Lululemon Athletica) and they appear to be again driven by differences in split adjustments.

Finally, I analyze discrepancies in the blue observations with differing coupons. The bottom panel of Figure A.3 zooms in on these observations and reveals beam-shaped patterns that appear to be caused by the platform applying shorter periods when calculating coupon payments. I manually check a number of these discrepancies and confirm with the prospectuses that the products pay fixed coupons independent of the performance of the underlying and that the coupon payment period used in the translated formulas is correct. To sum up, neither of the closer inspections reveal systematic errors in the translated payoff formulas or in the calculated returns.

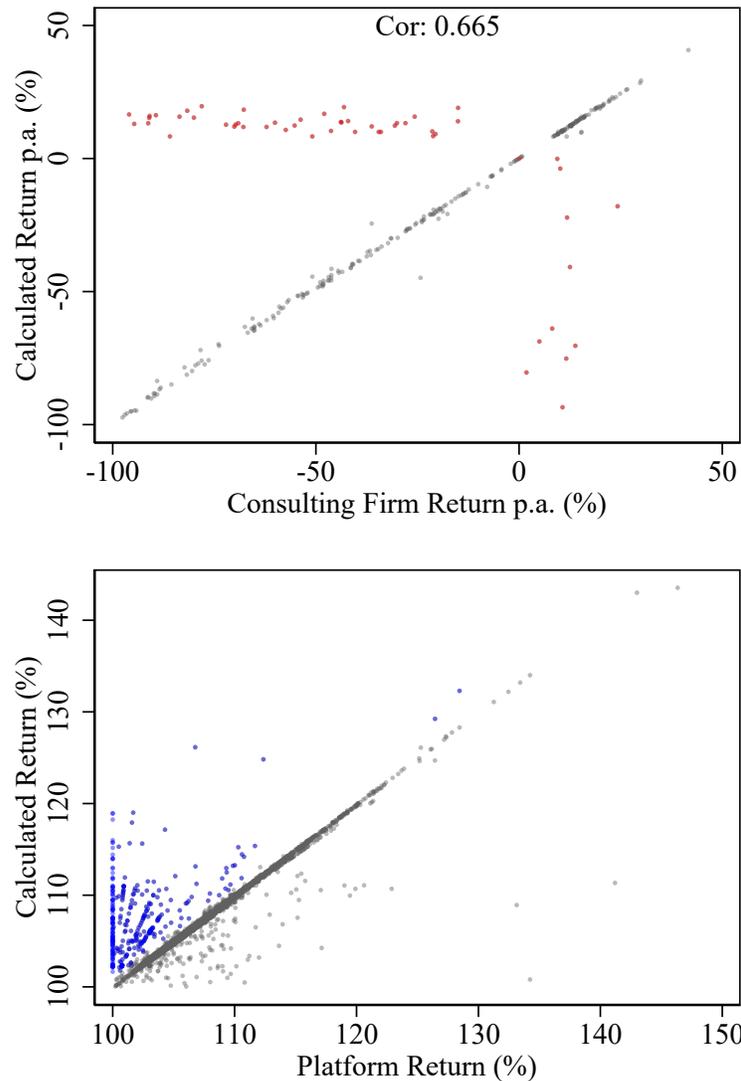


Figure A.3: Inspection of Discrepancies

The top figure plots returns calculated from translated payoff formulas against annualized return published by a consulting firm for those cases in which the returns from translated formulas disagree with the platform returns about repayment of full principal at maturity ($N = 298$). The bottom figure is the zoomed top-right quadrant of returns calculated from translated payoff formulas against total returns reported by the platform.

B Valuation Procedure

This section describes the valuation model used to derive fair values of the products. I use this model to calculate product fees derived from their fair values at issuance, as well as secondary prices used to construct YEP indices and calculate product deltas for benchmarking.

I follow previous literature (C  lerier and Vall  e, 2017) and use a local volatility model to value the exotic products in my sample. This approach is a generalization of the Black Scholes model with constant volatility used, for example, by Henderson and Pearson (2011). The main advantage of the local volatility model is that it treats volatility as a deterministic function of spot level and time and, by construction, it correctly prices vanilla options for all input strikes and maturities. Although the typical products in my sample are not vanilla, they are equivalent to portfolios of vanilla options, and one can therefore derive their model-free fair values by static or quasi-static replication.³ Therefore, the local volatility model will also produce correct and approximately model-free values of these exotic payoffs.

I start by building the local volatility surface, $\sigma(t, S_t)$. The model requires a full continuum in strikes and maturities one wishes to price, but in practice, the researcher observes only a discrete set of option prices. To interpolate volatilities between these discrete values, I use Andreasen and Huge (2011) interpolation, which is non-parametric and arbitrage-free.⁴ As the discrete set of market-implied volatilities, I use the out-of-the-money part of volatility surface from OptionMetrics.

To extrapolate volatilities for strikes and maturities not spanned by the OptionMetrics volatility surface, I use the volatility of the nearest maturity or strike price. In my context, the effect of extrapolation will be unimportant because the majority of the products are well spanned by the OptionMetrics volatility surface. The surfaces span up to two-year maturities, which is longer than the maximum possible maturity for 95% of the products. The average strike at one-year maturity and delta of -20 is 76% of the initial underlying price, whereas the average product barrier is 74%. In addition, in cases in which the product barrier is not spanned by the volatility surface of OptionMetrics, extrapolating from the nearest strike is a conservative assumption because the products have a negative vega and their volatility surfaces have a negative skew.

Because most of the products are linked to a single-name stock, I extend the local volatility diffusion with a negative jump on ex-dividend dates. The amount of the absolute dividend payments is extrapolated from the past 12-month history.

Whereas the underlying diffusion model is the same for all products in my sample, the

³See, e.g., Allen (2013) for static replication of digital options and quasi-static replication of barrier options.

⁴LexiFi Apropos, the pricing software that I use for validation and C  lerier and Vall  e (2017) use for their fair value estimation, uses the same local volatility interpolation.

valuation approach I use depends on the product payoff type. This approach is computationally efficient, which is important because I need to estimate fair values for 4.5 million product-day combinations. When possible, I use the finite difference method to price the options embedded in the product payoff. That is, I first decompose the product into a fixed-income component and option component and then use the model described above to price the embedded option.⁵ I price embedded binary options using a vanilla call spread, where the price of the vanilla options is calculated with a finite difference scheme.

For path-dependent autocallable products, I use Monte Carlo simulations. To reduce variance, I use a deterministic ("quasi-random") Sobol sequence generator and simulate 10,000 paths over 252 days in a year.

I use a similar valuation model to calculate the expected returns of YEPs in Section 3.4. The main difference is that to quantify expected returns, I consider the undiscounted payoffs of YEPs and use the expected return on the underlying, μ_t , instead of the risk-free rate r_t , in Eq. 5. I then again use either static replication, finite difference method, or Monte Carlo simulations to quantify the expected returns depending on the YEP payoff type. This estimation slightly underestimates the expected total returns of YEPs, because it ignores potential reinvestment of coupons paid before the product maturity.

B.1 Validation

I use commercial pricing software to cross-validate the accuracy of this complex estimation procedure. I thank LexiFi for providing me with a temporary license of their pricing software to perform this task. At the time when the license was available, my sample consisted of 21,390 products that I use for validation. The sample covered all payoff types and therefore allows me to validate accuracy across all payoffs as well as all valuation approaches.

⁵Tables D.1 and D.2 show examples of this decomposition.

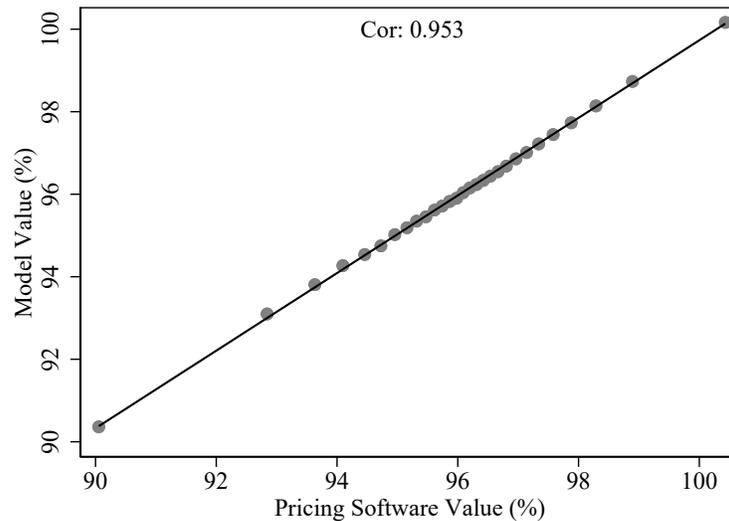


Figure B.1: Validation with Pricing Software

The figure plots a binned scatterplot and fitted line of product fair values at issuance estimated with the valuation model described in Section 3.1 against fair value estimates from pricing software. The sample consists of 21,390 products.

Figure B.1 shows a binned scatterplot of my estimates and estimates from the pricing software. The estimates are highly correlated ($\rho = 95\%$) and their means are nearly identical with a mean difference in fair values of 2.9 basis points. Both valuations use the same product data, valuation inputs, and valuation approach, and therefore any differences are driven only by a differences in the implementation. For example, the pricing software uses proprietary adjuster and control variate methods, which decompose the product into a statically replicated component and residual component priced with numerical methods. These adjusters improve precision of the estimation, which is important in a context where the precision of each estimate is crucial. In my context, the precision of my results emanates from the large sample size of more than 28,000 products, in which idiosyncratic discrepancies in fair value estimates cancel out.

As a second validation check, I compare my estimates with fair value estimates reported by issuers in the prospectuses. Figure B.2 plots binned scatterplots of fair values and annualized margins. Both estimates are again highly correlated, $\rho = 61\%$ for fair values, and $\rho = 87\%$ for annualized margins, although the correlation is lower than with the pricing software. One can expect the correlation with issuer estimates to be lower for many reasons. First, the value of the products is highly sensitive to the valuation inputs used, and these may differ for issuer estimates.⁶ For example, based on the instructions published by the SEC in 2013, issuers can

⁶For example, in Henderson and Pearson (2011) the correlation of baseline estimates with estimates under alternative assumptions on implied volatility is only 81 – 87%.

use their internal funding rate to price the fixed-income component of the products. To price the option component, issuers should mostly use mid-market inputs. These inputs should be broadly consistent with the OptionMetrics volatility surface I use, but differences may exist in the way issuers interpolate and extrapolate from observed market prices.

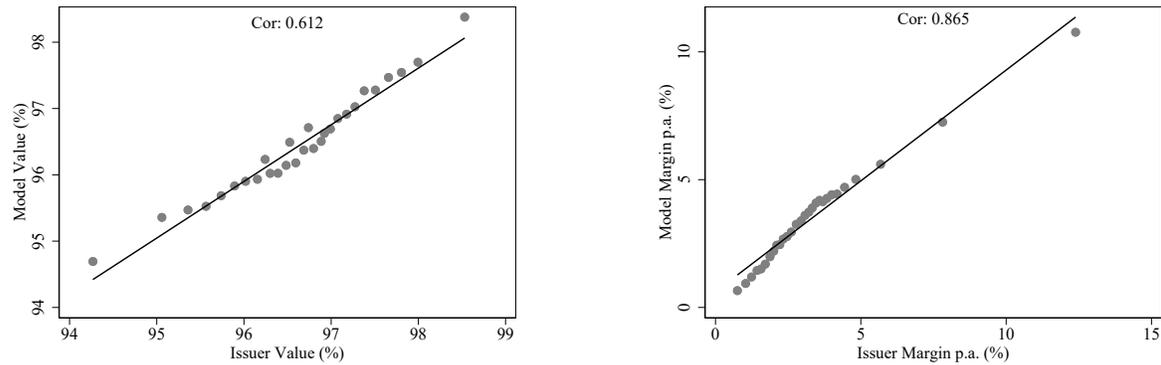


Figure B.2: Validation with Issuer Estimated Values

The figure plots a binned scatterplots and fitted lines of product fair values at issuance/issuer margins p.a. estimated with the valuation model described in Section 3.1 against issuer estimates disclosed in the prospectuses. The sample consists of 6,372 products with available issuer estimated values.

Second, issuers may use different valuation approaches. To estimate the values, issuers use proprietary pricing models and do not disclose any details on the models or assumptions used. Differences in valuation procedures are more important for more complex products that are more sensitive to the choice of valuation approach. I find that in the sample of less complex products that can be priced without Monte Carlo simulations, the correlation between mine and issuer fair values increases from 61% to 81%.

Finally, important issuer heterogeneity may be present in valuation inputs and valuation procedures that may be hard to capture by any consistent valuation method used by a researcher. Consistent with this hypothesis, I find the R^2 from regressing fair values (annualized margins) on issuer estimates increases from 0.375 (0.748) to 0.428 (0.775) and 0.489 (0.814) when adding issuer and issuer \times product type fixed effects, respectively.

For the purposes of my study, potential systematic bias is more important than idiosyncratic variation in product values. I find that neither comparison with the pricing software, nor with the issuer estimates shows economically important differences that could drive my results. An ultimate test of any systematic bias is comparison of the ex-ante estimated margins or fees with the ex-post abnormal performance. Section 3.5 shows my ex-post abnormal returns are in line with the ex-ante estimates—both when using delta-hedged abnormal returns as well as when estimating alphas from benchmark regressions of YEP indices.

As the last validity check, I again compare my estimates with secondary prices reported

in TRACE. This comparison serves both as a validation check of translating product payoffs presented in the previous section, as well as of the fair value estimation procedure. Errors in modeling payoffs should show up as significant non-linear deviations from the 45-degree line, whereas systematic differences in fair value estimates should show up as mean differences between mine and TRACE prices. Figure B.3 shows that my estimates are systematically higher than secondary prices. The figure plots binned scatterplot of my estimates and TRACE prices as well as the 45-degree line. Across the range of secondary prices, my estimates lie above the 45-degree line, which is consistent with brokers charging markdowns on secondary prices. The size of the markdown, discussed in more detail in Section C.1, is in line with markdowns reported in other bond markets.

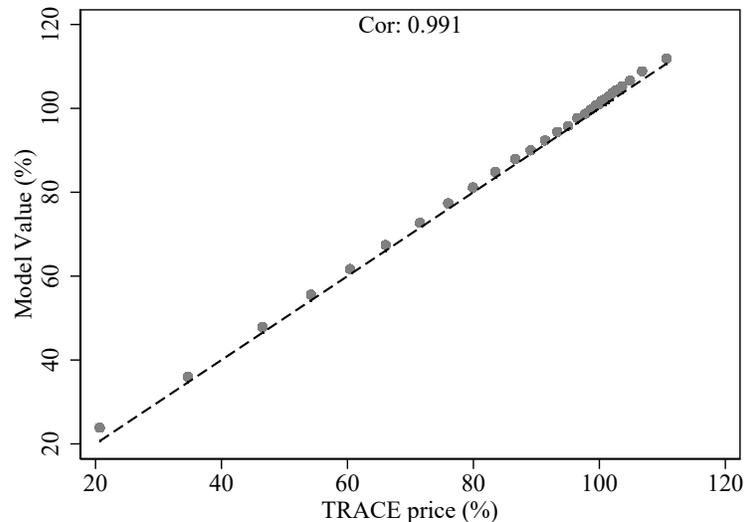


Figure B.3: Validation with TRACE Prices

The figure plots a binned scatterplot of product secondary fair values estimated with the valuation model described in Section 3.1 against secondary market prices reported in TRACE. The sample consists of 24,746 product-day-price observations. The dashed line is a 45-degree line.

B.2 Comparison with Analytic Valuation

My baseline valuation uses the local volatility model and therefore correctly prices options across the volatility smile. In a previous version of this paper, I assumed volatility to be constant. This assumption significantly simplifies the valuation procedure and allows me to use analytic closed form valuation to derive fair values for the majority of the products, similarly as in Egan (2019). In this section, I compare the results of the analytic valuation with the valuations based on local volatility.

For the purpose of analytic valuation, I assume the risk-neutral process for the underlying

asset is

$$dS = (r - q)Sdt + \sigma Sdz, \quad (1)$$

where S denotes the underlying asset price, r the risk-free rate of return, q the dividend yield provided by the stock, and σ the volatility of the stock. In other words, I assume constant volatility as well as continuous dividend yield. Under these simplifying assumptions, I can use text-book formulas to price all but autocallable products.

Products without a knock-out feature can be valued using standard textbook (e.g. Hull 2018) formulas for option valuation. The price of a *plain vanilla* put option with strike price K is calculated as

$$p = Ke^{-rT}N(-d_2) - S_0e^{-qT}N(-d_1), \quad (2)$$

where

$$d_1 = \frac{\ln(S_0/K) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln(S_0/K) + (r - q - \sigma^2/2)T}{\sigma\sqrt{T}}.$$

The *down-and-in* put option with barrier H is valued as

$$p_{di} = -S_0N(-x_1)e^{-qT} + Ke^{-rT}N(-x_1 + \sigma\sqrt{T}) + S_0e^{-qT}(H/S_0)^{2\lambda}[N(y) - N(y_1)] \\ - Ke^{-rT}(H/S_0)^{2\lambda-2}[N(y - \sigma\sqrt{T}) - N(y_1 - \sigma\sqrt{T})],$$

where

$$\lambda = \frac{r - q + \sigma^2/2}{\sigma^2}$$

$$x_1 = \frac{\ln(S_0/H)}{\sigma\sqrt{T}} + \lambda\sigma\sqrt{T}$$

$$y = \frac{\ln(H^2/(S_0K))}{\sigma\sqrt{T}} + \lambda\sigma\sqrt{T}$$

$$y_1 = \frac{\ln(H/S_0)}{\sigma\sqrt{T}} + \lambda\sigma\sqrt{T}.$$

The price for an *asset-or-nothing* option equals to

$$p_a = S_0e^{-qT}N(-d_1). \quad (3)$$

Cash-or-nothing call and put options are priced as

$$c_c = Qe^{-rT}N(d_2) \quad (4)$$

and

$$p_c = Qe^{-rT}N(-d_2). \quad (5)$$

As valuation inputs, I use bi-linearly interpolated implied volatility from the four options with the closest expiry dates before and after the option expiry date and the closest strike prices above and below the option strike price. In cases in which one or more of the four options are not available, I follow [Henderson and Pearson \(2011\)](#) and take the implied volatility of the option with the closest expiry date and the nearest strike price. As the risk-free rate, I use the linearly interpolated rate from the two OIS rates with the nearest maturities.

To estimate the dividend yield of stocks, I follow the methodology of OptionMetrics. I consider the dividend yield to be constant and equal to the most recent dividend payment divided by the most recent closing price. Unless a dividend payment date is already declared, I project the ex-dividend dates by extrapolating from the past dates and the most recent dividend payment frequency. The predicted dates extend up to the maximum maturity of a product. I then sum all predicted dividend payments over the life of a product and convert them to a continuous dividend yield.

Table [B.2](#) compares results from this analytic valuation with the results from local volatility valuation. The margins from analytic valuation are higher by 63 (9) basis points using the equal-weighted (volume-weighted) average. To understand the drivers of the difference, note the products have a negative vega, the underlying options have a negative skew, and their strike prices are often out of the money. Therefore, analytic valuation based on the interpolated volatility around the option strike price will tend to undervalue the products, because it applies higher constant volatility across all spot prices. This effect will be more pronounced for products with more skewed volatility, such as for short-term products. That said, given that the mean difference on a volume-weighted basis is not economically meaningful and that the analytic valuation is computationally lighter and more transparent, it may serve as a reasonable valuation procedure in many research applications.

Table B.1: Margins from Analytic Valuation

The table reports estimates of product margins obtained from analytic valuation described in this section and from the local volatility valuation model described in Section 3.1. The sample consists of 20,139 products with available closed form valuation under the assumption of constant volatility. Standard errors are clustered at the issuer level and reported in brackets.

Margins		
	Analytic	Local Volatility
Mean	4.21 (0.23)	3.58 (0.23)
Vol.-wtd. Average	3.76	3.67
p25	2.88	2.26
p50	4.16	3.50
p75	5.44	4.70
Observations	20,139	20,139

C Secondary Market Markdowns

As a supplementary evidence about the ex-post performance of YEPs, I estimate the markdowns investors incur when selling the products before maturity. The products have relatively short maturities and are intended to be held until maturity; therefore, selling them before maturity should be less common. Nevertheless, I observe 24,746 sales from retail investors back to the product providers using TRACE data.⁷

Panel A of Table C.1 presents summary statistics of secondary prices from TRACE and respective fair values for the same product-day combinations. Consistent with brokers charging markdowns on the products, the average secondary price is lower than the fair value estimate from my valuation model. I define the secondary market markdown as the difference between the secondary price and fair value divided by the fair value.

Panel B of Table C.1 shows the estimated markdowns categorized based on the time from issuance or time from maturity. In addition, Figure C.1 plots a binned scatterplot and kernel densities for the first 300 days after issuance and the last 300 days before maturity. Consistent with the previous literature and SEC findings,⁸ I find issuers provide price support in the few months after issuance. Immediately following the issue date, the secondary price is on average 1% higher than the product fair value. I also find the markdowns decline as the products approach maturity. The average markdown is 2% in the period of more than 90 days after issuance and more than 30 days until maturity, but less than 1% in the last 30 days.

These markdowns add to the poor performance of the YEP market and are consistent with the relatively high trading costs in other bond markets (Bessembinder, Spatt, and Venkataraman, 2020).

⁷I follow the method of Dick-Nielsen (2014) to clean TRACE data.

⁸See SEC letter available at <https://www.sec.gov/Archives/edgar/data/895421/000000000013009967/filename1.pdf>

Table C.1: Secondary Market Markdowns

Panel A reports summary statistics of secondary prices. *Fair value* is estimated using the pricing model described in Section 3.1. *Secondary price* (in %) is the price reported in TRACE divided by the primary price reported in TRACE. Panel B reports estimated markdowns (in %) calculated as Secondary price/Fair value -1 . Standard errors are clustered at the issuer level and reported in brackets.

Panel A: Summary statistics of secondary prices

	Mean	S.D.	p1	p25	p75	p99	Observations
Fair value	86.7	22.6	20.6	75.9	102.5	113.1	24,746
Secondary price	85.3	22.6	18.7	75.0	101.0	110.7	24,746

Panel B: Markdowns

	< 90 days after issuance	> 90 days after issuance and 30 days until maturity	< 30 days until maturity
Mean	-0.71 (0.33)	-2.00 (0.49)	-0.95 (0.37)
p25	-1.83	-3.18	-2.41
p50	-0.62	-1.70	-1.09
p75	0.60	-0.39	0.04
Observations	3,228	19,738	1,784

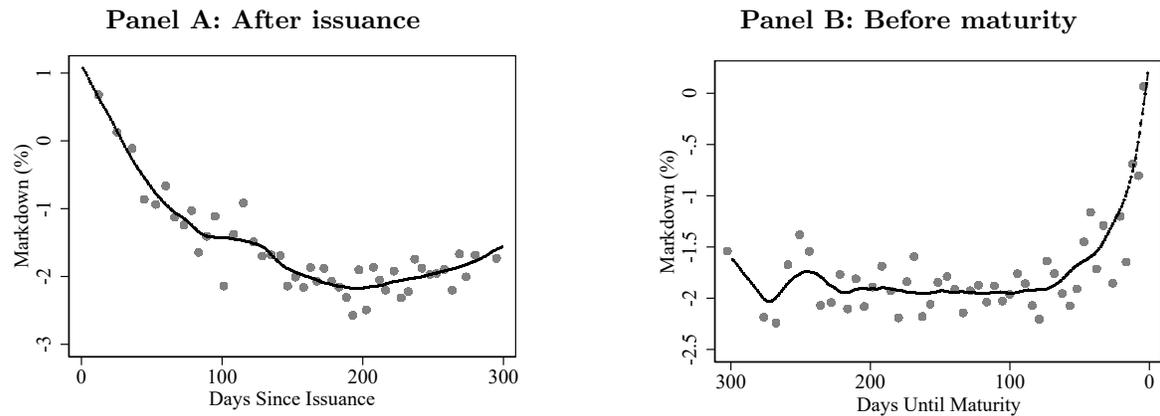


Figure C.1: Secondary Market Markdown

The figures plot binned scatterplots and kernel densities of markdowns in secondary prices. Panel A covers the first 300 calendar days after issuance and excludes the last 90 days until maturity. Panel B covers the last 300 days before maturity and excludes the first 90 days after issuance. The secondary price markup is the percentage discount in the price reported to TRACE over the fair value estimated using the pricing model described in Section 3.1.

D.1 Additional Figures and Tables

UBS AG Trigger Phoenix Autocallable Optimization Securities

UBS AG \$132,000.00 Securities Linked to the common stock of Cornerstone OnDemand, Inc. due on February 11, 2016

Final Terms											
Issuer	UBS AG, London Branch										
Principal Amount	\$10.00 per security. The Securities are offered at a minimum investment of 100 Securities at \$10.00 per Security (representing a \$1,000 investment) and integral multiples of \$10.00 in excess thereof.										
Term	Approximately 12 months, unless called earlier.										
Underlying Equity	The common stock of Cornerstone OnDemand, Inc. If the closing price of the underlying equity is equal to or greater than the coupon barrier on any observation date, UBS will pay you the contingent coupon applicable to such observation date.										
Contingent Coupon	If the closing price of the underlying equity is less than the coupon barrier on any observation date, the contingent coupon applicable to such observation date will not be payable and UBS will not make any payment to you on the relevant coupon payment date. The contingent coupon will be a fixed amount based upon equal quarterly installments at the per annum contingent coupon rate. Contingent coupons are not guaranteed and UBS will not pay you the contingent coupon for any observation date on which the closing price of the underlying equity is less than the coupon barrier. The table below reflects the contingent coupon rate of 13.24% per annum. Amounts in the table below may have been rounded for ease of analysis.										
	<table border="1"> <thead> <tr> <th>Observation Date*</th> <th>Contingent Coupon (per security)</th> </tr> </thead> <tbody> <tr> <td>04-May-2015</td> <td>\$0.3310</td> </tr> <tr> <td>04-Aug-2015</td> <td>\$0.3310</td> </tr> <tr> <td>04-Nov-2015</td> <td>\$0.3310</td> </tr> <tr> <td>04-Feb-2016</td> <td>\$0.3310</td> </tr> </tbody> </table>	Observation Date*	Contingent Coupon (per security)	04-May-2015	\$0.3310	04-Aug-2015	\$0.3310	04-Nov-2015	\$0.3310	04-Feb-2016	\$0.3310
Observation Date*	Contingent Coupon (per security)										
04-May-2015	\$0.3310										
04-Aug-2015	\$0.3310										
04-Nov-2015	\$0.3310										
04-Feb-2016	\$0.3310										
Contingent Coupon Rate	*Observation dates are subject to the market disruption event provisions set forth in the Trigger Phoenix Autocallable Optimization Securities product supplement ("TPAOS product supplement"). 13.24% per annum (or approximately 3.310% per outstanding quarter).										
Automatic Call Feature	The Securities will be called automatically if the closing price of the underlying equity on any observation date is equal to or greater than the initial price. If the Securities are called on any observation date, UBS will pay you on the corresponding coupon payment date a cash payment per Security equal to your principal amount plus the contingent coupon otherwise due on such date pursuant to the contingent coupon feature. No further amounts will be owed to you under the Securities.										
Payment at Maturity (per Security)	If the Securities are not called and the final price is equal to or greater than the trigger price and coupon barrier, UBS will pay you a cash payment per Security on the maturity date equal to your principal plus the contingent coupon otherwise due on the maturity date. If the Securities are not called and the final price is less than the trigger price, UBS will pay you a cash payment on the maturity date of significantly less than the principal amount, if anything, resulting in a loss of principal that is proportionate to the decline of the underlying equity, for an amount equal to $\$10 + (\$10 \times \text{underlying return})$.										
Underlying Return	$\frac{\text{Final Price} - \text{Initial Price}}{\text{Initial Price}}$										
Closing Price	On any trading day, the last reported sale price (or, in the case of NASDAQ, the official closing price) of the underlying equity during the principal trading session on the principal national securities exchange on which it is listed for trading, as determined by the calculation agent.										
Initial Price	\$32.54, which is the closing price of the underlying equity on the trade date. The initial price is subject to adjustments in the case of certain corporate events, as described in the TPAOS product supplement.										
Trigger Price/Coupon Barrier	\$22.78, which is 70.00% of the initial price of the underlying equity. The trigger price and coupon barrier are subject to adjustments in the case of certain corporate events, as described in the TPAOS product supplement.										
Final Price	The closing price of the underlying equity on the final valuation date. The final price is subject to adjustment in the case of certain corporate events, as described in the TPAOS product supplement.										
Trade Date	February 4, 2015										
Settlement Date	February 9, 2015										
Final Valuation Date	February 4, 2016 (subject to postponement in the event of a market disruption event, as described in the TPAOS product supplement)										
Maturity Date	February 11, 2016 (subject to postponement in the event of a market disruption event, as described in the TPAOS product supplement)										
Coupon Payment Dates	Five business days following each observation date, except the coupon payment date for the final valuation date will be the maturity date.										
CUSIP	90274L673										
ISIN	US90274L6737										
Valoren	26948314										

The estimated initial value of the Securities as of the trade date is \$9.25 for Securities linked to the underlying equity. The estimated initial value of the Securities was determined as of the close of the relevant markets on the date of this final terms supplement by reference to UBS' internal pricing models, inclusive of the internal funding rate. For more information about secondary market offers and the estimated initial value of the Securities, see "Key Risks - Fair value considerations" and "Key Risks - Limited or no secondary market and secondary market price considerations" in this final terms supplement.

Figure D.1: Pricing supplement with disclosed issuer estimated value

Investor Suitability

The suitability considerations identified below are not exhaustive. Whether or not the Notes are a suitable investment for you will depend on your individual circumstances, and you should reach an investment decision only after you and your investment, legal, tax, accounting and other advisors have carefully considered the suitability of an investment in the Notes in light of your particular circumstances. You should also review "Key Risks" on page 4 of this pricing supplement and "Risk Factors" on page 8 of the accompanying product supplement.

The Notes may be suitable for you if:

- You fully understand the risks inherent in an investment in the Notes, including the risk of loss of your entire initial investment.
- You can tolerate a loss of all or a substantial portion of your initial investment and are willing to make an investment that may have the full downside market risk of an investment in the Reference Underlying.
- You believe the Final Price of the Reference Underlying is not likely to be below the Conversion Price and, if it is, you can tolerate receiving shares of the Reference Underlying at maturity that are worth less than your initial investment or may have no value at all.
- You understand and accept that you will not participate in any appreciation in the price of the Reference Underlying and that your return at maturity is limited to the Coupon Payments.
- You are willing to accept the risks of owning equities in general and the Reference Underlying in particular.
- You can tolerate fluctuations in the price of the Notes prior to maturity that may be similar to or exceed the downside price fluctuations of the Reference Underlying.
- You are willing to invest in the Notes based on the Coupon Rate set forth on the cover of this pricing supplement.
- You are willing and able to hold the Notes to maturity, a term of approximately 6 months, and accept that there may be little or no secondary market for the Notes.
- You are willing to accept the credit risk associated with Deutsche Bank AG, as Issuer of the Notes, and understand that if Deutsche Bank AG defaults on its obligations you may not receive any amounts due to you, including any repayment of your initial investment at maturity.

The Notes may *not* be suitable for you if:

- You do not fully understand the risks inherent in an investment in the Notes, including the risk of loss of your entire initial investment.
- You require an investment designed to provide a full return of your initial investment at maturity.
- You are not willing to make an investment that may have the full downside market risk of an investment in the Reference Underlying.
- You believe the Final Price of the Reference Underlying is likely to be below the Conversion Price, which could result in a total loss of your initial investment.
- You cannot tolerate receiving shares of the Reference Underlying at maturity that are worth less than your initial investment or may have no value at all.
- You seek an investment that participates in the full appreciation in the price of the Reference Underlying or that has unlimited return potential.
- You are not willing to accept the risks of owning equities in general and the Reference Underlying in particular.
- You cannot tolerate fluctuations in the price of the Notes prior to maturity that may be similar to or exceed the downside price fluctuations of the Reference Underlying.
- You are not willing to invest in the Notes based on the Coupon Rate set forth on the cover of this pricing supplement.
- You are unable or unwilling to hold the Notes to maturity, a term of approximately 6 months, and seek an investment for which there will be an active secondary market.
- You are not willing or are unable to assume the credit risk associated with Deutsche Bank AG, as Issuer of the Notes for all payments on the Notes, including any repayment of your initial investment at maturity.

Figure D.2: Investor suitability section of product prospectus

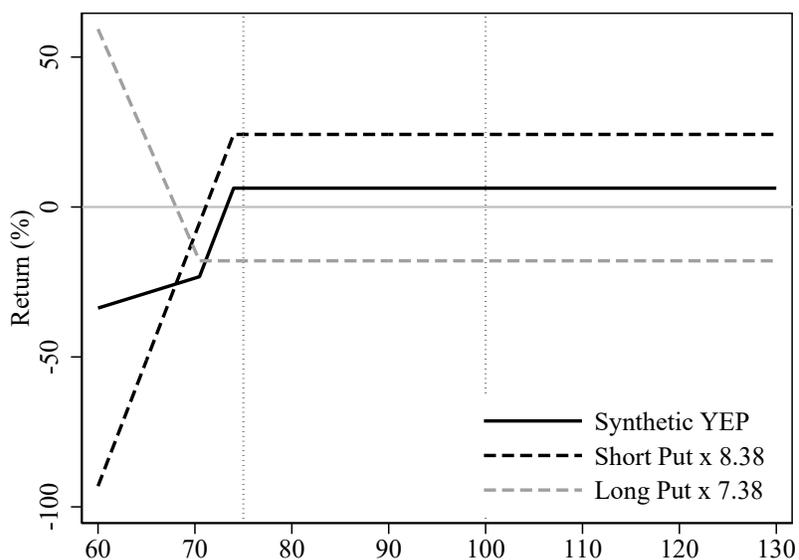


Figure D.3: Construction of Synthetic YEP

The figure shows the construction of synthetic YEP for the example product described in Fig. 1 and Table 1. Total return is on the y -axis. Prices of the underlying on the x -axis are normalized to 100 at issuance. The synthetic YEP can be thought of as a combination of a short position in a put option and bull put spread.

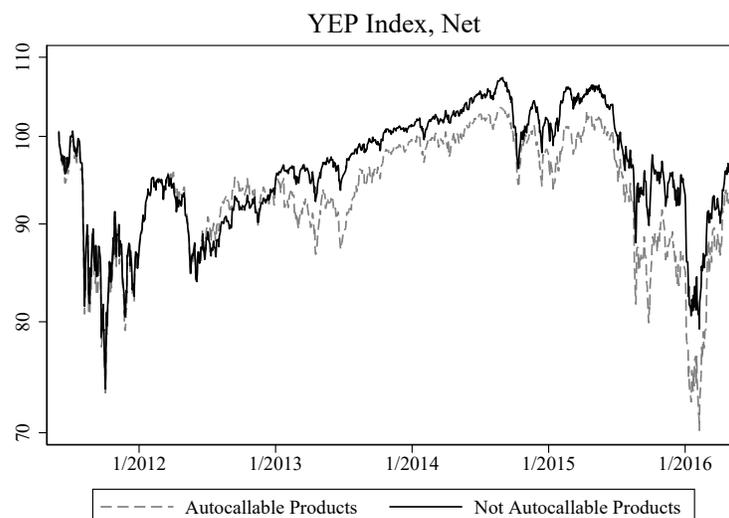


Figure D.4: YEP Net Index for Autocallable and Not-Autocallable Products

The figure shows decomposition of YEP Net Index to autocallable and not autocallable products. The figure plots the log index growth between June 2011 and May 2016 when there are at least 100 outstanding autocallable and not-autocallable products on each day. The construction of the YEP indexes is described in Section 5.1. The figure demonstrates the contrast between manipulation-proof YEP indexes and biased performance measures based on average annualized returns of products with early terminations conditional on performance as described in Section 5. The average total and annualized return of autocallable products issued after June 2011 and maturing before June 2016 are -0.32% and 6.88% ($N = 5,962$), respectively. The average total and annualized return of not-autocallable products issued after June 2011 and maturing before June 2016 are 0.01% and 0.13% ($N = 6,364$), respectively. Although autocallable products appear to perform significantly better in terms of annualized returns, the YEP Net Index shows their average performance was in fact worse than for not-autocallable products over the sample period.

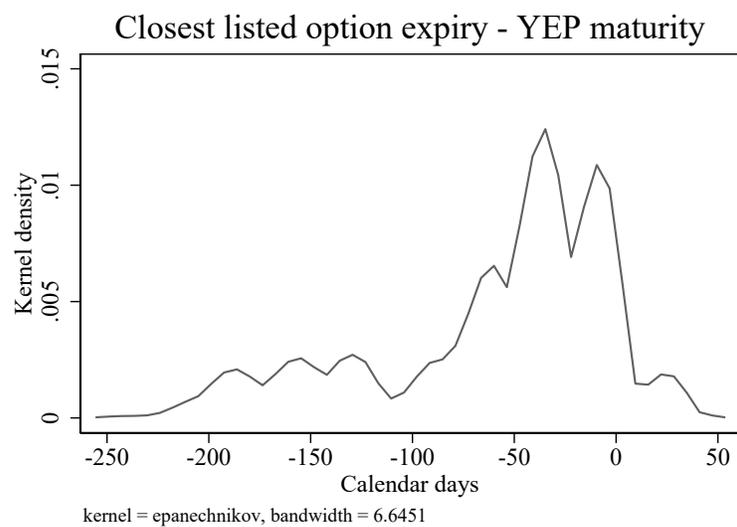


Figure D.5: Difference between YEP Maturity Dates and Option Expiry Dates

The figure plots the kernel density of the difference in calendar days between YEP maturity and expiry date of the option with the closest expiry date to YEP maturity. The sample consists of 17,000 YEPs used in the baseline specification of Table 6.

Table D.1: Example of Plain Vanilla Product

The table presents an example of a product with no exotic feature, its payoff description, translation into a mathematical formula, and decomposition into a bond and an option. Prices of the underlying are normalized to 100 at issuance.

Panel A: Product characteristics

Name	Airbag Yield Optimization Notes
Issuer	UBS
CUSIP	90272G254
Volume	\$0.35 million
Year	2014
Term	12 months
Coupon rate	5.22%
Underlying	Apple

Panel B: Payoff translation and decomposition

Payoff description This is an income product linked to the share of Apple. The product offers a coupon of 5.22% p.a., paid monthly throughout the investment period. At maturity, the product offers a capital return of 100%, if the final share level does not fall by more than 5% from its initial level. Otherwise the capital return equals 100%, minus 1.0526% for every 1% fall in excess of the initial 5% fall, paid in shares.

Translated formula

$$P_T = \begin{cases} 100 & \text{if } S_T \geq 95 \\ 100 - 1.0526(95 - S_T) & \text{otherwise} \end{cases}$$

$$P_m = 5.22 \times \Delta t \quad \text{for monthly observation dates } m = 1, \dots, M$$

Decomposition

Long bond, 5.22% monthly coupon
 Short 1.0526× put, $K = 95$

Table D.2: Example of Product with Binary Feature

The table presents an example of a product with a binary feature, its payoff description, translation into a mathematical formula, and decomposition into a series of conditional options. Prices of the underlying are normalized to 100 at issuance.

Panel A: Product characteristics

Name	Trigger Yield Optimization Notes
Issuer	Barclays Bank
CUSIP	06741K361
Volume	\$7.21 million
Year	2011
Term	6 months
Coupon rate	9.93%
Underlying	Nabors Industries

Panel B: Payoff translation and decomposition

Payoff description This is an income product linked to the share of Nabors Industries. The product offers a coupon of 9.93% p.a., paid monthly throughout the investment period. At maturity, the product offers a capital return of 100%, if the final share level does not fall by more than 25% from its initial level. Otherwise the capital return equals 100%, minus 1% for every 1% fall, paid in shares.

Translated formula

$$P_T = \begin{cases} 100 & \text{if } S_T \geq 75 \\ S_T & \text{otherwise} \end{cases}$$

$$P_m = 9.93 \times \Delta t \quad \text{for monthly observation dates } m = 1, \dots, M$$

Decomposition

Long bond, 9.93% monthly coupon
Short cash-or-nothing put, $K = 75, Q = 100$
Long asset-or-nothing put, $K = 75$

Table D.3: Example of Product with Knock-Out Feature

The table presents an example of a product that can terminate early, its payoff description, translation into a mathematical formula, and decomposition into a series of conditional options. Prices of the underlying are normalized to 100 at issuance. *Maximum term* is the maximum maturity if the product does not terminate early. *Expected term* is estimated under the risk-neutral probabilities of termination on each observation date. *Effective term* is calculated by evaluating the conditions for early termination on each observation date.

Panel A: Product characteristics

Name	Trigger Phoenix Autocallable Notes
Issuer	RBC
CUSIP	78010UZA8
Volume	\$0.5 million
Year	2014
Maximum term	18 months
Expected term	5 months
Effective term	3 months
Coupon rate	13.8%
Underlying	Facebook

Panel B: Payoff translation

Payoff description This is an income product linked to the share of Facebook. The product offers a coupon of 13.8% p.a., paid quarterly, if the share level does not fall by 30% or more from its initial level on the applicable quarterly observation date. Otherwise, no coupon is paid for that observation period. The product can terminate early on any quarterly observation date if the share level is greater than or equal to its initial level. In that case, the product terminates with a payout equal to 100% of the capital plus the coupon. At maturity, the product offers a capital return of 100%, if the final share level does not fall by 30% or more from its initial level. Otherwise the capital return equals 100%, minus 1% for every 1% fall, paid in cash or shares.

Translated formula

$$P_T = \begin{cases} 100 & \text{if } S_n < 100, \forall n = 1, \dots, M-1 \text{ and } S_T > 70 \\ S_T & \text{if } S_n < 100, \forall n = 1, \dots, M-1 \text{ and } S_T \leq 70 \\ 0 & \text{otherwise} \end{cases}$$

$$P_m = \begin{cases} 100 + 13.8 \times \Delta t & \text{if } S_n < 100, \forall n = 1, \dots, m-1 \text{ and } S_m \geq 100 \\ 13.8 \times \Delta t & \text{if } S_n < 100, \forall n = 1, \dots, m-1, S_m > 70 \text{ and } S_m < 100 \\ 0 & \text{otherwise} \end{cases}$$

Table D.4: Regressions Explaining Dominated Products

The table reports coefficient estimates and t -statistics from regressions of indicator variables for dominated product (Column 1 and 2), excess synthetic coupon rate (Column 3 and 4), and issuance volume (Column 5 and 6) on product characteristics. The sample covers 14,383 products that can be statically approximated with up to two positions in put options and lending and for which data on commissions are available. *Dominated* is an indicator variable for YEPs dominated by synthetic counterparts constructed in Section 4. *Excess coupon* is the difference between the implied coupon rate of synthetic counterpart and YEP coupon. *Volume* is product issuance volume in million \$. *Commission* is broker's commission in %. *Margin* is product margin estimated in Section 3.3 minus broker's commission. Columns 1–4 and 6 include product type fixed effects where the omitted category includes products with embedded vanilla options, and *Barrier* and *Digital* product types embed barrier and binary options, respectively. t -statistics based on standard errors clustered at the issuer level are reported in parentheses. *, **, *** indicate statistical significance at the 5%, 1%, and 0.1% level, respectively.

Dominated products, volume, and product characteristics

			Excess		Volume	
	Dominated	Dominated	coupon	coupon	Volume	Volume
Dominated					0.0449 (0.253)	0.0243 (0.607)
Commission	0.124*** (6.595)	0.117*** (4.773)	0.0173*** (5.773)	0.0118** (4.340)		0.0974* (2.517)
Margin	0.0896*** (10.193)	0.0918*** (15.849)	0.0126*** (9.731)	0.0126*** (8.630)		-0.0859* (-2.936)
Coupon		-2.137*** (-8.375)		-0.422*** (-5.800)		1.337 (1.473)
Barrier		2.578*** (8.280)		0.314*** (9.798)		-0.968** (-3.760)
<i>Product type</i>						
Barrier	-0.608*** (-11.730)	-0.522*** (-5.970)	-0.060*** (-9.008)	-0.0247* (-2.948)		-2.023 (-1.864)
Digital	-0.450*** (-9.413)	-0.359*** (-12.862)	-0.045*** (-9.906)	-0.0183** (-4.336)		-0.487** (-3.174)
<i>Underlying × Maturity ×</i>						
Issuer × Month FE	No	Yes	No	Yes	No	Yes
Observations	14,383	4,054	14,383	4,054	14,383	4,054
R^2	0.188	0.729	0.257	0.837	0.000	0.752

Table D.5: Regressions Explaining Underlying Selection

The table reports coefficient estimates and t -statistics from a linear probability model explaining underlying selection in YEPs. The sample covers month-stock observations between January 2006 and September 2015 for non-financial U.S. stocks that belong to the 750 largest stocks by market capitalization at the end of the previous calendar year. The dependent variable is an indicator variable equal to 1 if the stock has been used as an underlying in the issuance month for at least one YEP. *Implied volatility* is the volatility of calls with 30-days maturity and delta = 50 from the volatility surface table of OptionMetrics at the end of the month preceding the issuance month. *Log market capitalization* is the natural logarithm of the stock market capitalization at the end of the calendar year preceding the issuance month. *Log 1-month volume* is the natural logarithm of the stock volume during the month preceding the issuance month. *3-month return* and *12-month return* are measured over 3 and 12 months preceding the issuance month. Columns 1–6 include month fixed effects. In addition, Column 6 includes underlying fixed effects. *** indicates statistical significance at the 0.1% level.

Probabilistic regression coefficients and t -statistics

Implied volatility	0.497*** (60.051)				0.483*** (53.187)	0.0645*** (5.919)
Log market capitalization		0.0622*** (61.675)			0.0360*** (27.056)	0.0649*** (22.422)
Log 1-month volume			0.0947*** (108.761)		0.0676*** (57.914)	0.0635*** (28.436)
3-month return				-0.0342*** (-4.450)	0.0374*** (5.236)	0.0120 (1.863)
12-month return				0.0448*** (15.394)	0.0484*** (18.043)	0.0271*** (10.762)
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Underlying FE	No	No	No	No	No	Yes

References

- Allen, Steven, 2013, *Financial Risk Management. A Practitioner's Guide to Managing Market and Credit Risk, 2nd edition*. John Wiley & Sons, Inc.
- Andreasen, Jesper, and Brian N. Huge, 2011, Volatility Interpolation. *Risk Magazine*, 24, 86–89.
- Bessembinder, Hendrik, Chester Spatt, and Kumar Venkataraman, 2020, A survey of the microstructure of fixed-income markets. *Journal of Financial and Quantitative Analysis*, 55, 1–45.
- Célérier, Claire, and Boris Vallée, 2017, Catering to investors through security design: Headline rate and complexity. *Quarterly Journal of Economics*, 132, 1469–1508.
- Dick-Nielsen, Jens, 2014, How to clean Enhanced TRACE data, Working paper.
- Egan, Mark, 2019, Brokers versus retail investors: Conflicting interests and dominated products. *Journal of Finance*, 74, 1217–1260.
- 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*, 100, 227–247.
- Hull, John, 2018, *Options, futures, and other derivatives, 10th Edition*. Pearson.