NBER WORKING PAPER SERIES

AN ANALYSIS OF THE PERFORMANCE OF TARGET DATE FUNDS

John B. Shoven
Daniel B. Walton

Working Paper 27971
http://www.nber.org/papers/w27971

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2020

Shoven serves as an Independent Trustee of the American Century Funds. He also was a member of the board of directors of Financial Engines from January, 2010 to July, 2018. The views expressed in this paper do not necessarily reflect the views of American Century Funds or Financial Engines (now Edelman Financial Engines), nor those of the National Bureau of Economic Research.

An Analysis of the Performance of Target Date Funds
John B. Shoven and Daniel B. Walton
NBER Working Paper No. 27971
October 2020
JEL No. D14,D91,G11,G41,J32

ABSTRACT

This paper presents a thorough evaluation of target date funds for the period 2010-2020. These funds have grown enormously in assets, reaching $1.4 trillion by the end of 2019. They account for approximately 24 percent of all of the assets in 401(k) accounts. The paper reports on the results of a style analysis evaluation of TDFs which results in their effective asset allocation. It examines the constant in the style analysis regressions which reflects the over- or under-performance of the funds relative to a passive benchmark with the same asset allocation. Lower cost TDFs tend to match the benchmark, whereas higher cost TDFs deviate considerably from theirs. We examine how TDFs performed in the stock market crash between 2/19/20 and 3/23/20. In this five week period, broad market averages fell by about one-third. We find that the value of long-dated TDFs (those with a target date of 2045 and beyond) fell by between 30 and 35 percent, whereas the 2025 funds, designed for people roughly 60 years old, lost between 20 and 25 percent of their value. We find that past performance only weakly influences future expected performance. As with equity funds in general in this period, TDFs with actively managed ingredient funds, on average, trailed the performance of their cheaper passively managed counterparts.

John B. Shoven
Department of Economics
Stanford University
Landau Economics Building
579 Jane Stanford Way
Stanford, CA 94305
and NBER
shoven@stanford.edu

Daniel B. Walton
Department of Economics
Stanford University
Landau Economics Building
579 Jane Stanford Way
Stanford, CA 94305
dwalton@stanford.edu

1. Introduction

Target date funds have become enormously important assets in defined contribution retirement accounts. More than half of all 401(k) participants have at least some money in target date funds and approximately 24 percent of the total assets in 401(k) accounts are in this type of fund. The total amount invested in TDFs was approximately $1.4 trillion at the end of 2019.

The performance of target date funds was subject to considerable scrutiny in the Great Recession of 2008-09, leading to Congressional hearings and a report from the General Accountability Office (GAO, 2011) suggesting changes in disclosures about them. What prompted us to initiate this review of the performance of target date funds was the stock market crash in February and March of 2020. How did they do this time? This led us to take a broad look at target date funds and their performance over the past 5 to 10 years.

We create a database of all target date funds and balanced funds in existence in the U.S. in the 2010-2020 timeframe. We examine their expense ratios and find that the distribution is bimodal with roughly half of the money in TDFs facing expenses of less than twenty basis points and most of the other half facing fees between 50 and 70 basis points. Target date funds are a type of "fund of funds" with the ingredient funds being mostly different kinds of equity and fixed income funds. By and large, the lower fee funds are based on passive (indexed) ingredient funds and the higher expense funds are based on actively managed component funds.

We examine the "glide paths" of target date funds (their asset allocation over time) by examining their monthly returns between the beginning of 2015 and the end of 2019 and finding the best fit relative to the returns on 13 factor ETF funds. This is an exercise in style analysis as introduced by William Sharpe (1992). In general, we find that the style analysis reveals that target date funds change their effective asset allocation in a manner consistent with their advertised glide paths.

We examine the constant in the style analysis regressions. This constant reveals whether the TDF had a higher or lower average monthly return than the best fitting set of reference or factor funds. It is a close relative to the CAPM alpha of the target date funds. The target date funds with relatively low cost, those with expense ratios of less than 30 basis points, have style analysis constants close to zero. This indicates that they perform very similarly to their best fitting set of factor funds. On the other hand, the higher cost TDFs have a relatively disperse distribution of style analysis constants. On average, the more expensive funds have lower and negative constants, but they are overrepresented amongst the worst and best performing TDFs. All 35 of the funds with a constant of negative 10 basis points per month or lower are higher expense TDFs.

We look at how target date funds performed during the stock market crash between February 19 and March 23, 2020. The broad U.S. stock market fell by about one-third in this five week period. We consider this period a kind of "natural stress test" for target date funds. Most of the long-horizon funds (those with a target date of 2045 and beyond) lost between 30 and 35 percent of their value. That is, they did not do significantly better than a pure equity fund. More interestingly, most of the short-horizon funds, such as the 2025 funds, lost between 20 and 25 percent of their value over the five weeks. While this didn't surprise us after having done the style analysis, it may have disappointed participants anticipating retirement in approximately five years. While the glide path does move to more conservative asset allocations as retirement approaches, for many funds it never becomes what could be considered low risk or safe. On average, the lower cost funds did slightly better than the higher cost funds, but the more impressive difference was that there was a significantly wider range of outcomes among the higher cost funds.

We examine the Sharpe ratio, the ratio of the average excess return to risk, of the funds for the 2015-19 period. Once again, the pattern is that the average low-cost fund had a superior Sharpe ratio compared to the higher cost funds. Finally, we examine the connection between returns in the 2010-14 period and the subsequent returns in the 2015-19 period. We find a weak positive connection on average between performances in the two periods. A one percent per year better performance in the first period predicts 9 basis points per year better performance in the second period. In other words, there is a large reversion to the mean. The implication of this is that even a high cost fund with a very strong prior record of returns has a lower expected future rate of return than a low cost fund.

We end the paper with a discussion of whether a moderate risk balanced fund with some guidance on how to choose amongst a family of balanced funds might present a better default option than a target date fund. TDFs have a one size fits all aspect to them where the only difference between employees is their age and hence target retirement date. People differ in many other important dimensions. A set of balanced funds (also called target risk funds) might allow people to better sort themselves out in terms of their risk tolerance and hence optimal asset allocation.

2. Data

The data used in the analysis come from several sources to get as comprehensive a view as possible of the universe of target date funds and balanced funds. We scraped both Morningstar and Yahoo! Finance databases for all funds with a "target date" or "balanced allocation" label, coming up with a total of 4,116 unique tickers, with 1,547 of those tickers bearing the label of "balanced funds" and 2,569 of the tickers being labeled as "target date funds." Historical returns, net asset value (NAV), and assets under management (AUM) come from the WRDS CRSP database of mutual funds, while current prospectus data comes from Morningstar, Yahoo! Finance, and Google Finance. Historical data runs from January 1, 2010 to April 30, 2020, and returns and NAV are recorded on a daily basis, while AUM is reported monthly.

An important feature of mutual funds is the various share classes available to the investor. Many funds have several share classes, and some offer 10 or more different types of shares. Importantly, share classes within a fund differ in fees and expenses charged to shareholders, as well as in other dimensions such as services offered or administration support to the sponsor. A typical retail investor will likely not have access to all share classes of a given fund due to minimum investment requirements. Share classes of mutual funds can be roughly divided into two categories: retail and institutional. Retail share classes generally have higher fees and expenses but feature low minimum required balances, while institutional share classes usually have lower fees and expenses but high minimum required investments. Consequently, retail share classes are more accessible to individual investors, while the high minimum investment requirements restrict institutional shares to investments made by institutions. It is important to note, however, that shares of retirement-oriented funds available to an individual through their employer such as a 401(k) or 403(b) account will often be institutional-type shares, since the employer will have many employees' accounts invested in the same fund, thus meeting the required minimum investment for an institutional share class.

Much of our analysis examines the relationship between the expense ratio and historical returns of a fund, so we seek a way to aggregate the expenses across share classes of a fund into a single representative number that describes how "expensive" any given fund is, without referring to a specific share class. Such an aggregation allows comparison of fees across funds. This aggregation is achieved by taking a weighted average of expense ratios across share classes within a fund on a monthly basis, where the weights are proportional to the AUM of each share class[1]. This weighting permits us to calculate the average expense charged on a dollar invested in the fund, across all investments in the fund. This means that, all else equal, a fund with more dollars in retail shares will have a higher average expense ratio than a fund with more dollars in institutional class shares. One can calculate an average weighted return as well in the same manner, either by taking a weighted average of the given returns of the various share classes, or by taking any share class return, adding back the expenses taken out over the interval of the return, and subtracting the average expense ratio calculated above.

[1] AUM for share classes are given on a monthly basis in the CRSP historical data, while expense ratios for share classes are given on an annual basis, as determined in the annual prospectus in the CRSP historical data. However, the historical expense ratio data for many funds is incorrect or incomplete, so we use the current expense ratio for each share class and fund and assume it is constant over time. This is reasonable, since historical data on expense ratios for share classes show they generally vary by 5 basis or less points over time.

Once we take a weighted average of the share classes, we are left with data at the fund level. This aggregate data contains 977 distinct funds, 612 of which are categorized as target date funds, and 365 of which are balanced funds. Table 1 shows a more detailed count of funds by category. Target date funds come in 5-year vintages, with the year being the expected date of retirement, or closest date to expected retirement. Balanced funds, on the other hand, are distinguished by the percentage of fund assets invested in equities; the general idea is that funds with higher equity exposure will have higher risk and higher expected returns.

Fund Category		# Of Funds
Balanced Funds		**365**
	Allocation--15% to 30% Equity	37
	Allocation--30% to 50% Equity	100
	Allocation--50% to 70% Equity	137
	Allocation--70% to 85% Equity	56
	Allocation--85%+ Equity	35
Target Date Funds		**612**
	Target-Date Retirement	69
	Target-Date 2015	33
	Target-Date 2020	55
	Target-Date 2025	56
	Target-Date 2030	57
	Target-Date 2035	55
	Target-Date 2040	56
	Target-Date 2045	55
	Target-Date 2050	56
	Target-Date 2055	55
	Target-Date 2060+	65

Table 1: Number of Funds of Each Category in the Aggregate Data

3. The Growth of Target Date Mutual Funds

The growth of TDF mutual funds between 2000 and 2020 was nothing short of astounding. Figure 1 shows the assets under management (AUM) data assembled by the Investment

Company Institute (ICI, 2020). The bar for 2000 is so short as to be hard to see, but the total assets at the end of that year in target date funds amounted to $8 billion. At the end of 2019, the assets had grown to $1,395 billion, representing an average compound annual growth rate of more than 31 percent. About 80 percent of TDF assets are held in defined countribution retirement accounts, with another 15 percent held in Individual Retirement Accounts (IRAs), leaving only 5 percent in other types of accounts.

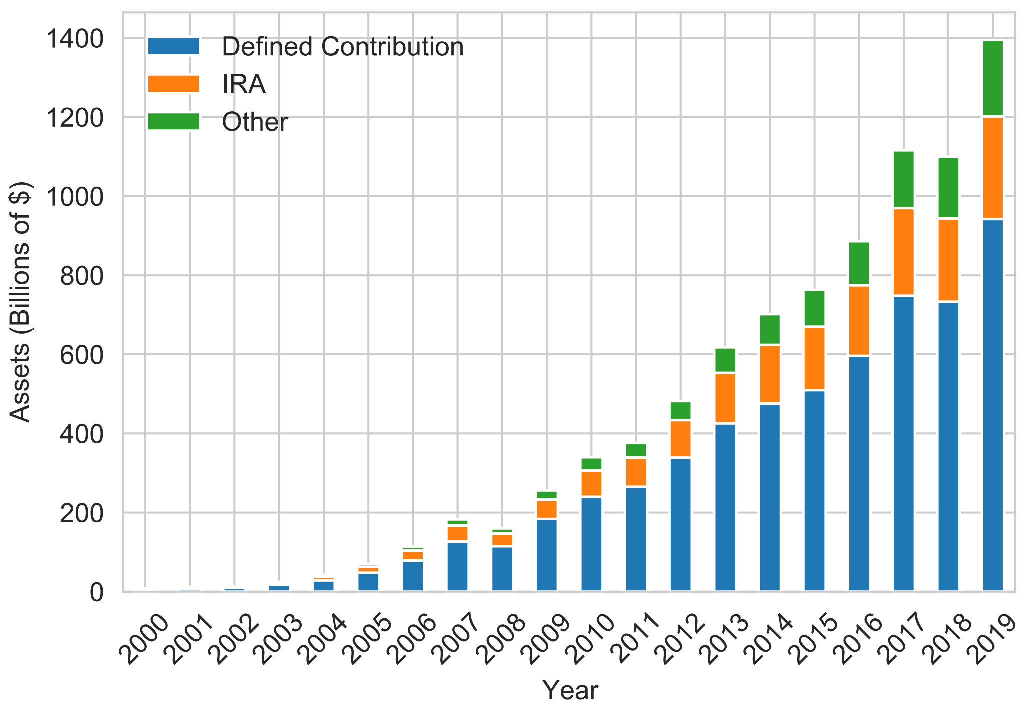

Figure 1: AUM in Target Date Funds by year

A significant contributor to this growth was the fact that in 2006 the Department of Labor ruled that target date funds could be a "Qualified Default Investment Alternative" (QDIA) in tax deferred retirement plans. The authority to designate QDIA status was part of the 2006 Pension Protection Act. This status provides plan sponsors some degree of protection from lawsuits resulting from poor investment performance if they use a TDF family as the default option in a plan with automatic enrolment. Balanced funds and professionally managed accounts also were given QDIA status, but target date funds proved to be far and away the most popular defaut option in such plans..

Figure 2 shows the growth of particular vintages of target date funds (such as 2025 or 2045 funds) over time. By 2019, both the 2025 and the 2030 funds had total assets which equaled the entire TDF assets only ten years earlier. Both ended the year with more than $200 billion in assets. It is not surprising that the funds with the most assets today are those whose target dates are relatively soon. The figure also hints at the fact that even the 2025 and 2030 funds took a big hit in the market crash of February and March 2020. The slowly growing but smooth line is the retirement class of target date funds and the one showing essentially no growth since 2015 is the 2015 target date fund category.

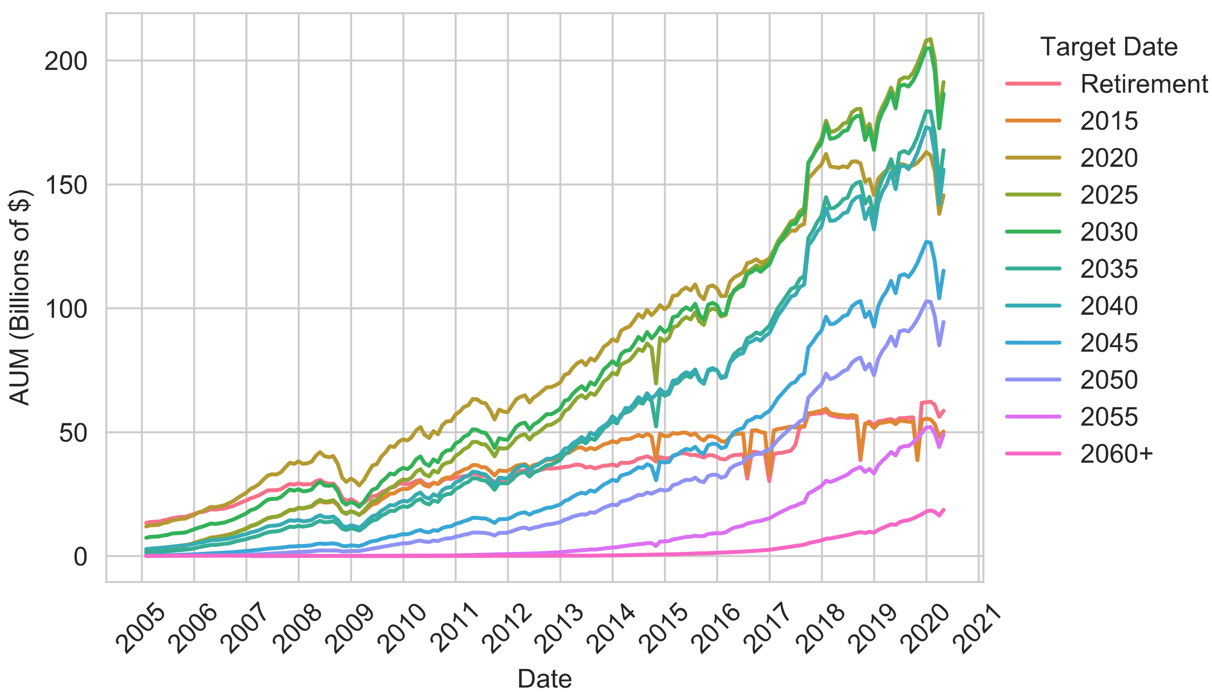

Figure 2: AUM in Target Date Funds by vintage over time

Target date funds can be thought of as an innovation on traditional balanced funds, which hold asset classes in roughly fixed proportions such as 60 percent equities and 40 percent bonds.

Target date funds are a kind of dynamic balanced fund which becomes more conservative as the target date approaches. Figure 3 shows the growth in assets under management over time for balanced funds. The biggest balanced funds are those whose equity allocations range from 50 to 70 percent. But even this most popular type of balanced fund has experienced a growth of AUM much slower than target date funds. In fact, the growth of assets in balanced funds has trailed their returns, indicating that they have not attracted positive net cash flows for the past five years. The extremely conservative and extremely aggressive funds have attracted very few assets over the past 20 years. The net flows of cash into balanced funds is shown in Figure 4.

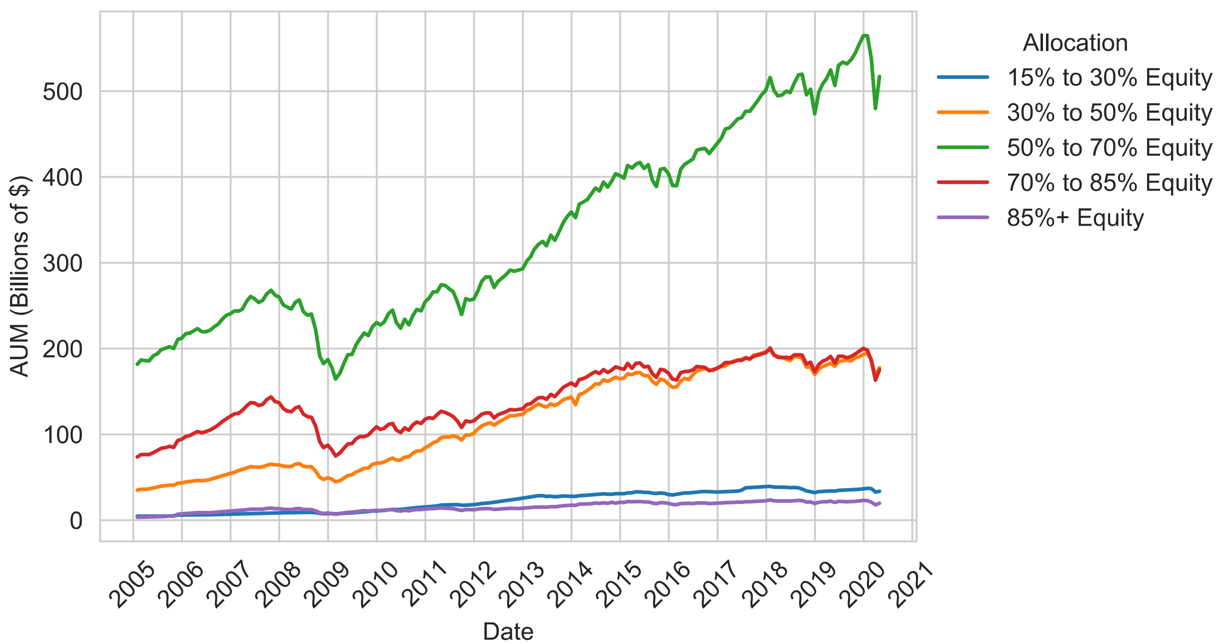

Figure 3: AUM in Balanced Funds by risk category over time

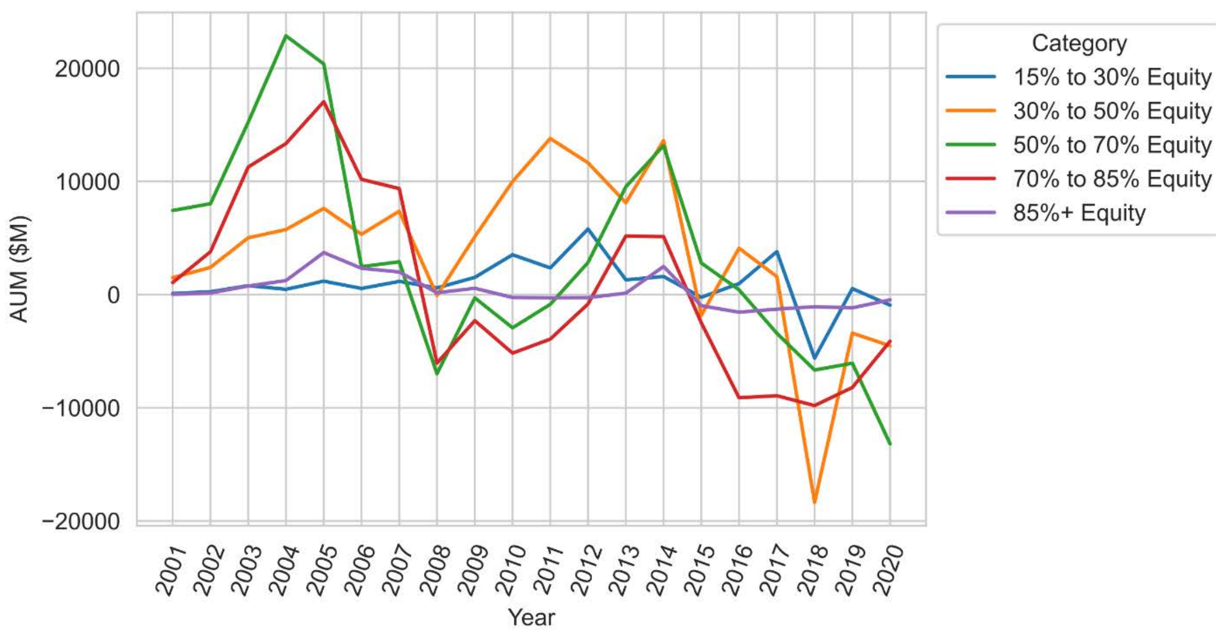

Figure 4: Cash Inflows into Balanced Funds, 2001-2020

The growth of target date funds has been impressive in other dimensions. For instance, by the end of 2016, fully 53 percent of all 401(k) participants had at least some money in TDFs. This figure is up from 19 percent at the end of 2006. Similarly, by the end of 2017, 23.8 percent of all 401(k) assets are held in TDFs, up from 3.1 percent in 2006. These 2016 and 2017 numbers are the latest available from the Investment Company Institute (2019), but the penetration of TDFs in 401(k) accounts has almost certainly continued since then. Most large 401(k) accounts offer TDFs. The figure was 82.4 percent in 2017. TDFs are even more prevalent in the accounts of young participants. In 2016, almost half (47.6%) of assets in accounts of participants in their 20s were invested in TDFs. People in their 60s, with their much larger accounts, held almost 20 percent (18.4% to be exact) in TDFs. It is too soon to say whether today's young people will continue to hold TDFs at a very high rate or whether they will change investment approaches as they age and their accounts grow.

These facts lead us to the fairly obvious conclusion that target date funds are large and important financial instruments. We feel that this justifies looking closely at their performance in recent years.

4. Expenses

Target date funds can be found with an extremely wide range of annual expense ratios, ranging from .08 percent to 1.98 percent per year. However, if you look at the assets under management in different expense ranges, you get the results shown in Figure 5. The distribtution is roughly bimodal, with nearly half of the total assets in target date funds imposing expense fees below 20 basis points (0.20%) with the bulk of the rest being subject to 50 to 70 basis points. Target date funds are a type of fund-of funds, where the ingredient funds are mostly various types of equity and fixed income mutual funds. The low cost TDFs are exclusively made up of passive (indexed) ingredient funds whereas the more expensive TDF funds have actively managed ingredient funds

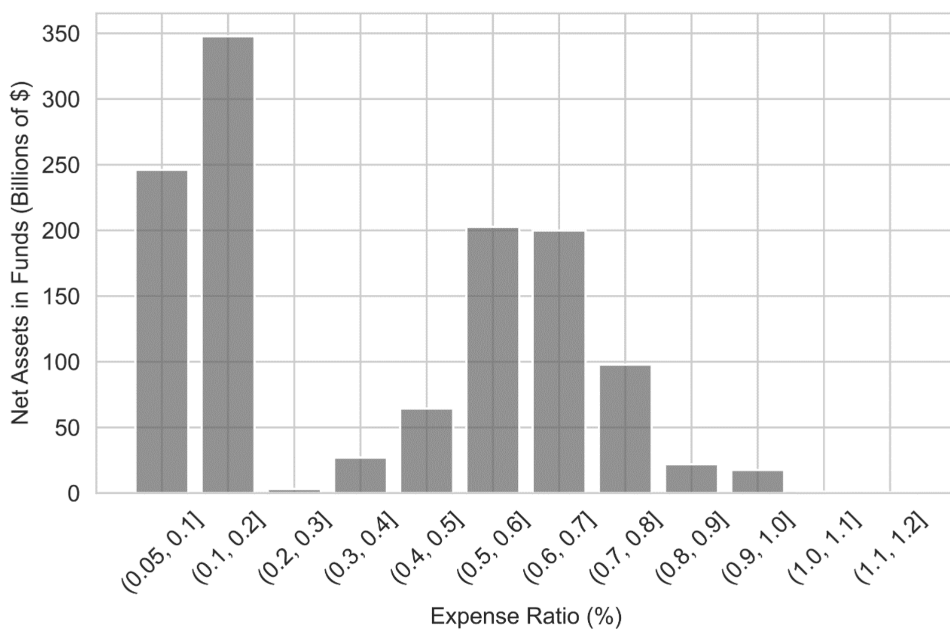

Figure 5: AUM in TDFs by Expense Ratio

Figure 6 displays much the same information in the form of a cumulative distribution function. It displays what fraction of the total assets in TDFs have expenses less than or equal to the expense ratios shown on the horizontal axis. In both figures the most striking fact is that almost no money in TDFs faces expense ratios between 20 and 40 basis points. The TDFs based on index funds are cheaper than that and those based on active portfolio management are more expensive than that. The separation in expense ratios is quite clean.

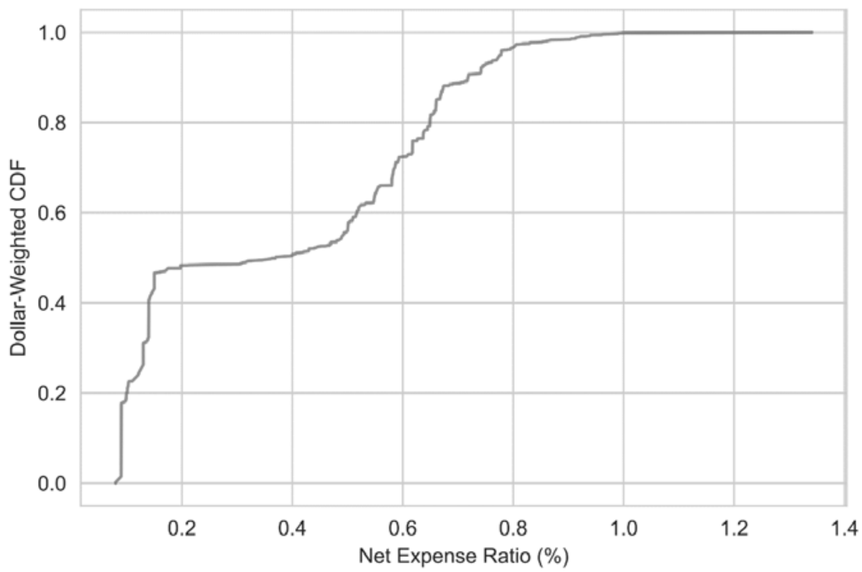

Figure 6: CDF of AUM in TDFs by Expense Ratio

Figure 7 shows the expense structure faced by investors in traditional balanced funds. It has a similar bimodal pattern to what Figure 5 shows for target date funds. This time there are very few dollars in balanced funds with expense charges between 20 and 50 basis points. The money in balanced funds facing expense charges between 80 and 120 basis points is much greater than with TDFs. We think that the explanation is that the bulk of TDF assets are in defined contribution plans which tend to be offered in the less expensive share classes (such as R classes or an institutional class). Balanced funds are more likely to be held in IRA accounts or taxable brokerage accounts and be held in more expensive retail share classes.

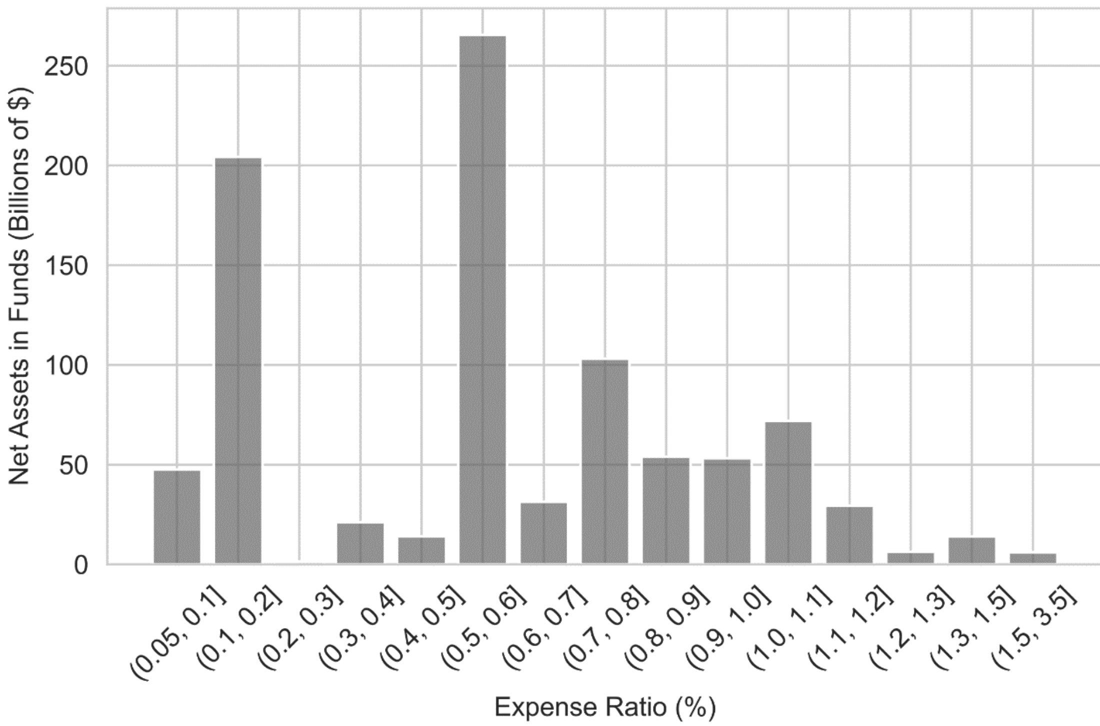

Figure 7: AUM by Expense Ratio for Balanced Funds

Figure 8 displays the same information as Figure 7, but this time in the form of a cumulative distribution function. Comparing Figures 6 and 8, one can see that the less expensive passive funds have a bigger share of the assets in target date funds than they do in balanced funds. With TDFS, almost half of the total assets face expense charges of 20 basis points or less. For balanced funds, less than 30 percent of the assets face such low fees. At the other end of the spectrum, only a very small percentage of the assets in TDFs face expense ratios of greater than 80 basis points, whereas with balanced funds, about 25 percent of the assets face expense fees of 80 basis points or more.

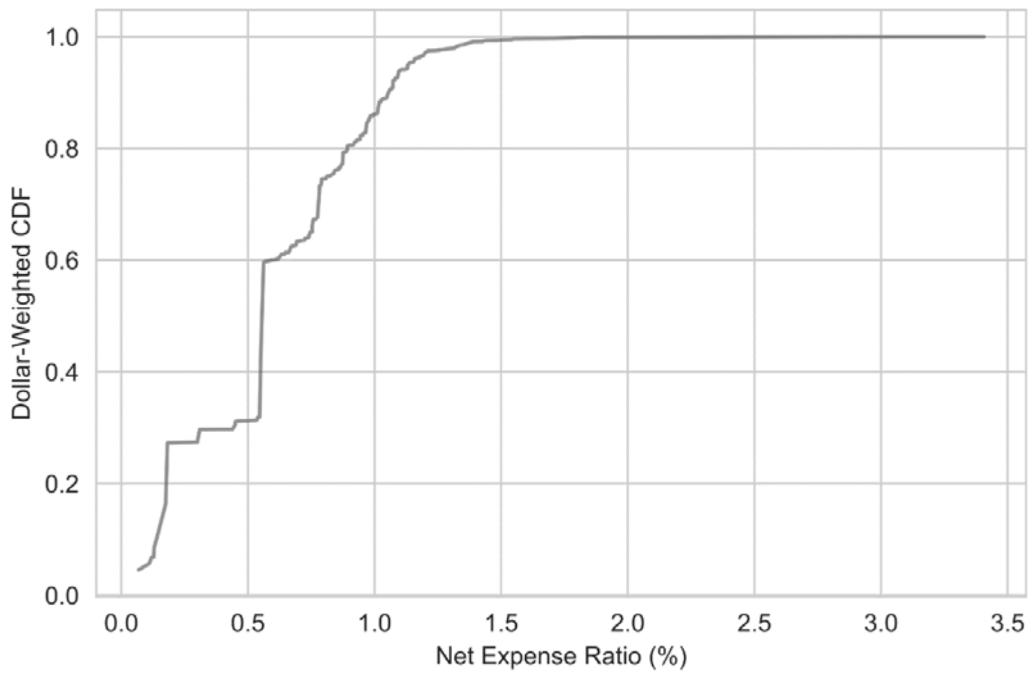

Figure 8: CDF of AUM in Balanced Funds by Expense Ratio

5. Style Analysis

To explain returns of a fund, it makes sense to determine the exposures of a fund's investments to various broad asset classes, such as large capitalization growth stocks, small-cap stocks, corporate bonds, government-issued debt, and so on. A fund may report its various holdings and one may use that data to determine exposures, but an interesting alternative is to simply use past returns of the fund and past returns of index funds that track the various broad asset classes in order to find what combination of asset classes best fit the past returns, thus determining the "effective" exposures of the fund's investments. Such a method was pioneered by William Sharpe (1992), and is often used in mutual fund analysis. We first describe the methodology of our style analysis before giving our results and discussing them.

Since style analysis determines a fund's exposure to a range of asset classes over time, we first need to select a period of data for which the investment style and asset allocation of the fund is relatively stable. This is important for target date funds, since they are advertised as changing exposure to risky asset classes over time, to decrease risk as the individual approaches retirement. On the other hand, to achieve the most accurate estimators, we would like to use a long period to have as many observations as possible. We chose the 5 year period beginning in January 2015 and ending in December 2019 to satisfy the tradeoff between more data and having a period of stable asset allocation. The target date funds themselves are labeled in 5-year intervals and are intended for any individual retiring in a 5-year window, indicating that it is likely for the fund to have similar asset allocations over a given 5 year period. Another subtler point in choosing data for style analysis consists of the return periods to use, where again there is a tradeoff; shorter return periods result in more data points, but tend to be noisier, having higher variance of annualized returns. We chose monthly returns to balance noise and observation frequency; and this appears to be a popular choice in the literature as well. Therefore, our data set used in the estimation of exposures consists of returns data on a monthly basis from January 2015 to December 2019, giving a total of 60 observations per fund. Note that this relies on the assumption that the fund exists for the entire period. This condition is not satisfied for every fund in the dataset, so we dropped all funds for which there wasn't a full time series of returns for the 5-year period. In particular, this had the effect of removing funds that were established since the beginning of 2015, which are mainly funds with a distant target date, such as the 2055, 2060, and 2065 funds.

We do not rule out the possibility of survivorship bias in the data as well, since we cannot tell how many funds that no longer exist are not included in our dataset, therefore our results concern the population of funds conditional on their survival through the analysis period. In total, the filter left us with 395 funds that reported returns during the 5-year period. We model each target date fund's return as a convex combination of the returns of its respective exposures, plus a constant, α, which is a measure of average performance of the fund relative to a portfolio that solely consists of diversified investments with the same exposures to the asset classes. We used 13 asset classes to explain fund exposures, using ETFs whose holdings are representative of the various asset classes. These ETFs and the asset classes they cover are given in Table 2. We refer to them as "factor funds" or reference funds.

Symbol	Exp Ratio	Asset Category	Annualized Return (01-2015 to 12-2019)
VUG	0.04%	U.S. LG CAP GROWTH	9.42%
VTV	0.04%	U.S. LG CAP VALUE	3.72%
IWR	0.19%	U.S. MID CAP	2.36%
VB	0.05%	U.S. SMALL CAP	1.30%
VGK	0.08%	EUROPE ALL CAP	-0.47%
VPL	0.08%	ASIA PAC ALL CAP	2.05%
BLV	0.05%	Long Term Bond	6.32%
BIV	0.05%	Int Term Bond	3.96%
BSV	0.05%	Short Term Bond	2.27%
BWX	0.35%	Foreign Sov Bond	0.63%
LQD	0.15%	U.S. Corp Bond	3.95%
VNQ	0.12%	REITs	1.34%
MBB	.06%	Mortgage Backed Securities	2.76%

Table 2: Reference ETFs (Factor Funds) Used in Style Analysis[2]

For the purposes of our analysis, we do not remove fund expenses from returns, so the interpretation of the constant α includes differences in the management costs of funds. Thus, returns are modeled as

$$R_{it} = \alpha_i + \sum_{j=1}^{13} \lambda_{ij} S_{jt} + \epsilon_{it}$$

where i indexes the fund whose exposures are to be determined, j indexes the asset class, and $t = 1, \ldots, T$ indexes the time of observation, with R_{it} the return of fund i, S_{jt} the return of asset class j, and ϵ_{it} an error term, for which we make the standard assumption of mean 0 and i.i.d. We impose the conditions that $\sum_{j=1}^{13} \lambda_{ij} = 1$ for all i, and $\lambda_{ij} \geq 0$. The former condition pins down one of the exposures, given the other 12. Therefore, we have 13 parameters (including

[2] Summary of Exchange-Traded Funds used to construct benchmarks in the style analysis: The first six ETFs track distinct, non-overlapping broad categories of equities, both in the United States and internationally. The next five ETFs capture non-overlapping categories of bonds and government-issued debt, both domestic and foreign. The final two ETFs cover alternative investments: real estate investment trusts and mortgage backed securities in the US.

α_i) to estimate. Define matrices $\beta_i = (\alpha_i, \lambda_{i1}, \dots, \lambda_{i13})'$, $R_i = (R_{i1}, \dots, R_{iT})'$, $X = \begin{pmatrix} 1 & S_{1,1} & \cdots & S_{13,1} \\ \vdots & \vdots & & \vdots \\ 1 & S_{1,T} & \cdots & S_{13,T} \end{pmatrix}$. We estimate the parameters by constrained least squares, with the minimization program formally expressed as

$$\min_{\beta_i}(R_i - X\beta_i)'\,(R_i - X\beta_i)$$

$$subject\ to$$

$$\iota_0\beta_i = 1, I_0\beta_i \geq 0$$

Where $\iota_0 = (0,1,\dots,1)$ of dimension 1x14, and $I_0 = \left(\vec{0}, I_{13}\right)$, a 14x13 matrix with the first column being a vector of zeros, and the other 13 columns making up the 13x13 identity matrix. This constrained least squares problem is a quadratic program, for which there exist many efficient numerical optimizers.

We ran the above constrained regression for all funds that existed during the entire Jan 2015-Dec 2019 period. Figure 9 shows the average effective asset exposures by three major classes (equities, bonds, and alternative investments), that are the sums of all the exposures within their respective categories, for each fund, and then aggregated by averaging over all funds within a given vintage. The "wicks" at the top of the bars in the graph display the 5th and 95th percentile allocation to the three major asset classes. The results are consistent with the publicized glide paths of target fund providers. The 2020 funds, on average, have an asset allocation of about 50 percent equities. The long duration TDFs have 85 to 90 percent equities. One may think of the different vintages (2020, 2025, 2030, etc.) as snapshots of how a single target date fund may choose effective asset allocations over its lifetime. For example, 5 years from now, we would expect a 2030 fund's exposures to look like a current 2025 fund's exposures. When retirement is far away, as in the 2045-2060+ target date funds, the funds have very high effective exposure to equities, over 80%, and correspondingly low exposure to bonds and alternative investments. As retirement becomes closer, the glide path to lower-risk investments sets in, and the exposure to equities decreases while exposures to bonds rise. While exposure to alternative investments is low for all vintages of funds, it is higher for funds closer to or already in retirement. From this evidence, we conclude that target date funds in effect are keeping their promise of following a glide path, moving to lower-risk and lower-expected return assets as the target retirement date approaches. Another way to empirically investigate the glide paths is to directly estimate expected return and volatility from past returns, rather than look at exposures to broad asset classes. We perform this type of analysis in Section 8.

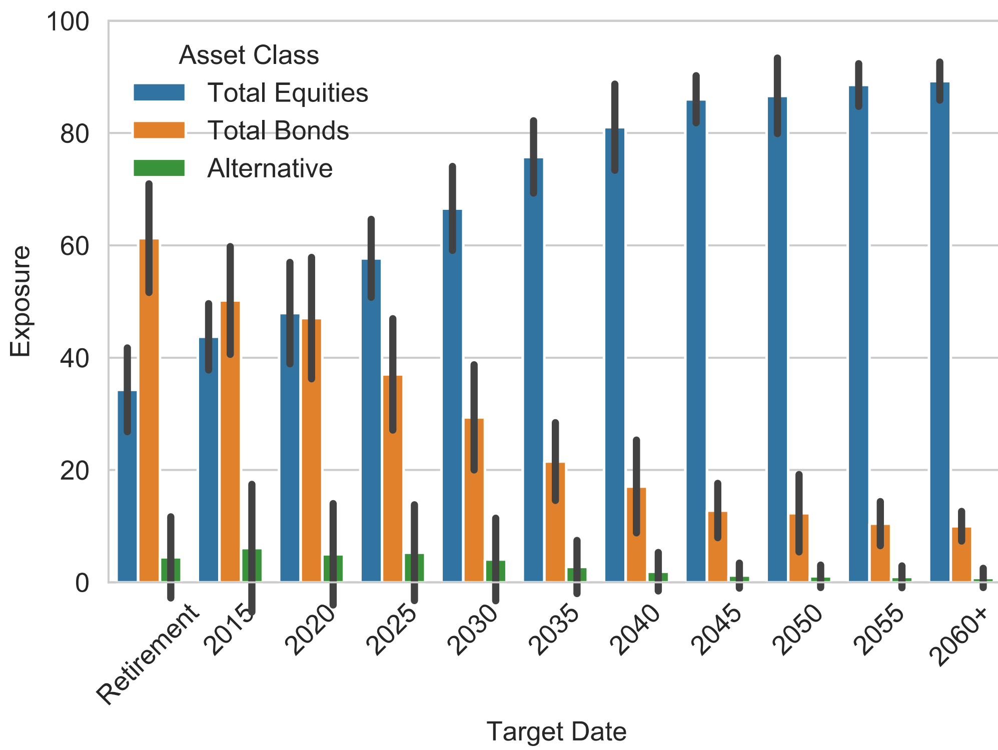

Figure 9: Exposures to asset classes from style analysis, target date funds

Figure 10 shows the style analysis results for traditional balanced funds. Once again, the results are not surprising. Effective equity exosure from the style analysis exactly corresponds to the risk category of the balanced funds.

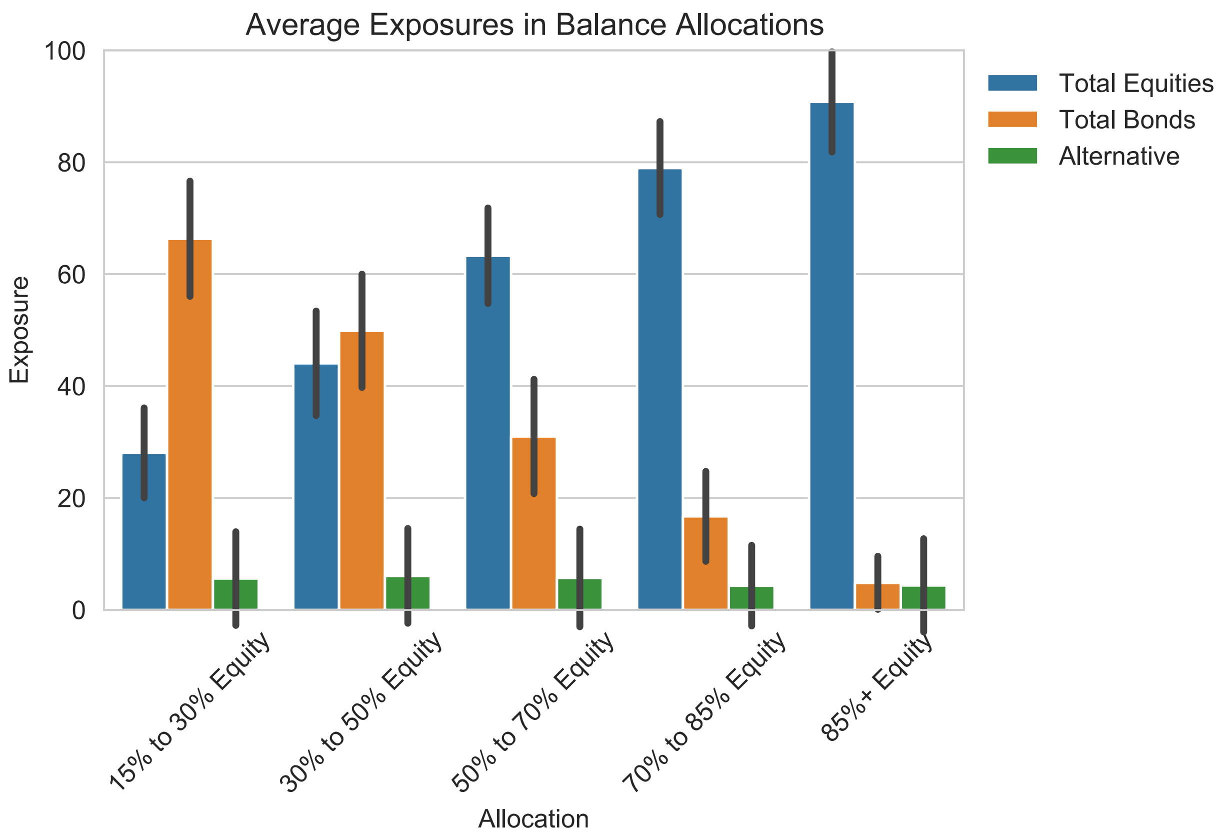

Figure 10: Exposures to asset classes from style analysis, balanced funds

6. A Natural Stress Test: February 19, 2020 to March 23, 2020

Even near term target date funds suffered large losses in the 2008-09 financial crisis, prompting Congressional hearings on their perceived poor performance. The market crash from February 19,2020 until March 23, 2020 offers a recent test of their performance in a dramatically bad equity market. In that five week period, the S&P 500 stock index fell 33.8 percent whereas the Vanguard Total Stock Market Index (VTI) fell 35 percent. The Vanguard Total Bond Market Index ETF (BND) had a total return of about -1.5 percent. Figure 11 shows how target date funds performed in the form of a "box and whiskers" graph. The vertical dimensions of the box shows the 75th and 25th percentile outcomes, the line in the middle of the box shows the median outcome and the end of the whishers show the 95th and 5th percentile outcomes. The points shown with diamound shapes are outliers. The figure indicates the most of the distant future TDFs (the 2045, 2050, 2055 and 2060+) lost between 30 and 35 percent of their value over this five week period, failing to do significantly better than an all equity portfolio invested in the S&P 500. Most of the near future TDFs (the 2025 funds) lost between 20 and 25 percent of their value. While the 2025 funds, presumably held by workers within roughly five years of retirement, did better than the long duration TDFs, they still lost more than one fifth of their value in five weeks which may have disappointed their shareholders given the promise of the glide paths becoming safer and more conservative as retirement approaches.

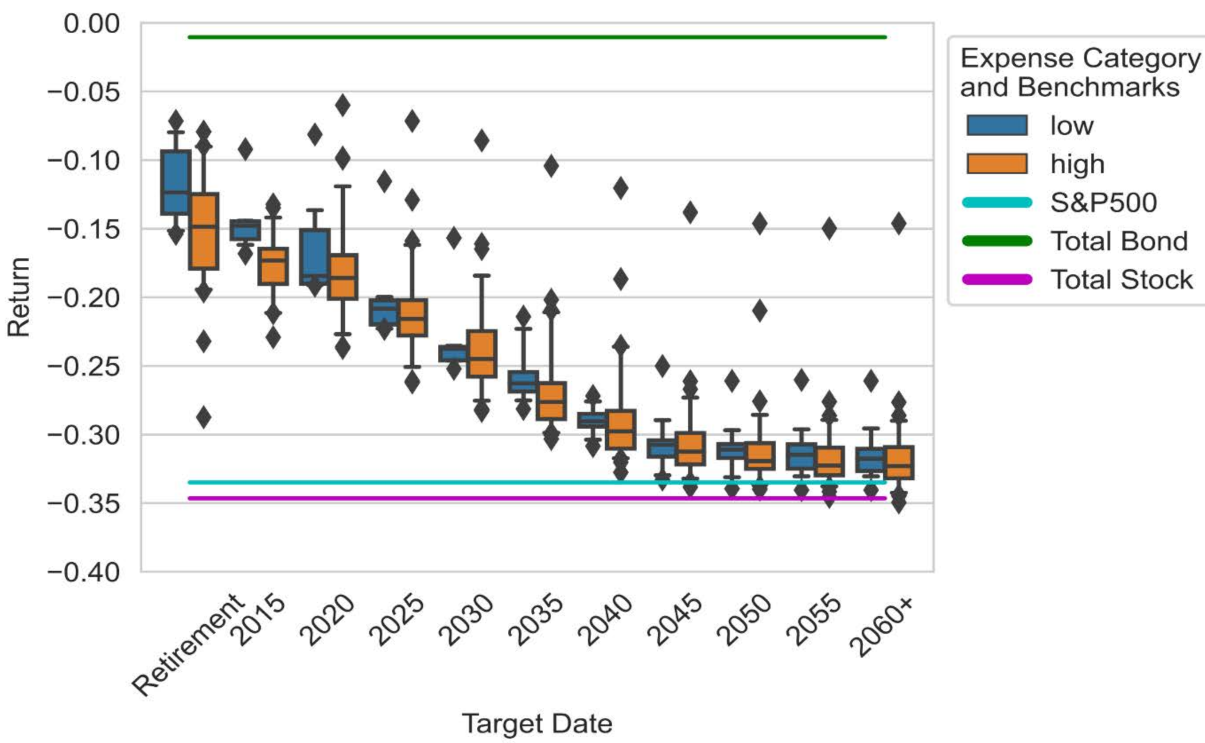

Figure 11: Returns for stress test period (Feb 19 to Mar 23, 2020, non-annualized), TDFs

Figure 11 shows other findings regarding returns during this stress test period. The figure divides target date funds into two groups in terms of expense ratios, those with fees less than 30 basis points (categorized as low expense funds) and those with fees over 30 basis points (the high expense category). Recall that few if any funds have expenses in the immediate neighborhood of 30 basis points. What is apparent is that the median performace is worse for the expensive funds than the low cost funds. However the range of outcomes is much wider for the high-cost funds. In all categories except the retirement one, the 95th percentile high-cost fund outperformed the 95th percentile low-cost fund. For several vintages, the 75th percentile high-cost fund outperformed the 75th percentile low-cost fund. For the poor performers in each category, the high-cost funds did worse than the low-cost funds. For instance, the 25th percentile high-cost fund did worse than the 25th percentile low-cost fund for all vintages. We think that the right interpretation of this is that the high-cost funds are almost all actively managed and deviate from the overall market in their portfolios. In some cases those deviations (or bets) pay off and in other cases they don't. The low-cost funds track the overall market much more closely.

We characterize target date funds as dynamic balanced funds. Figure 12 looks at how traditional balanced funds did in the stress test period, February 19, 2020 to March 23, 2020. The answers are not surprising. The balanced funds with 30 to 50 percent exposure to equities did approximately as well as the 2020 target date funds, whereas the balanced funds with 85+ percent equity exposure performed similarly to the 2045. 2050, 2055 and 2060+ target date funds. In terms of expense category, the balanced funds display exactly the same pattern as TDFs. In particular, the range of returns was much wider for the higher cost actively managed balanced funds than for the lower cost passive ones. We do note that while the 85+ percent equity balanced funds and the long duration TDFs performed similarly, the 85%+ balanced funds failed to attract assets in the marketplace.

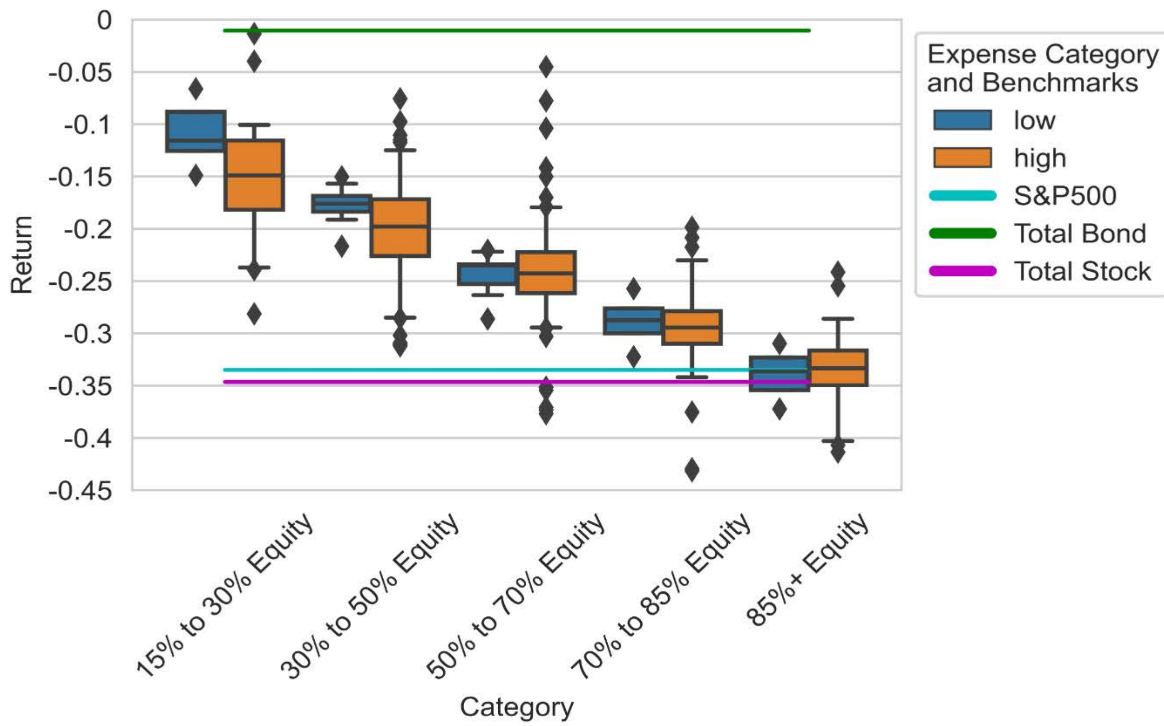

Figure 12: Returns for stress test period (Feb 19 to Mar 23, non-annualized), Balanced Funds

7. The Style Analysis Constant

The style analysis approach, as described above, serves a role beyond simply determining asset class exposures; it also creates a benchmark for each fund, against which we may measure a fund's performance. The regression constant α is equal to the difference $R_{it} - \sum_{j=1}^{13} \lambda_{ij} S_{jt} - \epsilon_{it}$, so in words, the estimate of α is the average deviation in returns of the fund and the returns of the portfolio of best fitting ETFs, since the errors average to 0. A positive constant indicates that the fund outperformed the best fitting set of reference ETFs, whereas a negative constant indicates underperformance relative to reference funds with the same asset allocation. This estimate, which we will call $\widehat{\alpha}$, thus measures the fund's performance relative to its asset allocation. The constant is closely related to the classic portfolio analysis measure also known as alpha, which is the difference between an asset's expected return and the expected return of a diversified portfolio of the same beta. Indeed, we expect that the portfolio of best fit from style analysis has a beta and expected return that lies on or very close to the capital market line, and we expect the actual fund and the benchmark portfolio to have the same beta, since the style analysis minimizes the variance between the fitted portfolio returns and the actual fund returns. To be clear, the returns used as an input to the style analysis regression are net of expenses; that is, the fund's net expense ratio is subtracted from the gross returns, for both the fund returns on the left hand side, as well as the ETF returns on the right hand side. This means that $\widehat{\alpha}$ is more relevant from an investor's perspective: it measures the net return performance of a fund relative to its benchmark portfolio, and for tax-advantaged accounts, this is essentially the take-home return.

In Figure 13, we plot the estimated constant $\widehat{\alpha}$ for each target-date fund against the expense ratio of the fund given at the beginning of the period. Since percentage returns were used in the style analysis estimations, the constant is in monthly percentage amounts, so, for instance, a fund with $\widehat{\alpha} = -0.05$ means that, on average for the 5-year period, the net return on the fund is 5 basis points per month lower than the return of the fitted portfolio. The corresponding annualized $\widehat{\alpha}$ would be equal to $100[(1 - 0.0005)^{12} - 1]$, which is approximately 60 basis points (slightly less), or 12 times the monthly value. We also plot (in red) the least squares fit from regressing the constant on an intercept and expense ratio, which is the equation $Constant = 0.012 - 0.079 * ExpenseRatio$.

We categorized funds with annual expense ratios greater than 30 basis points as "high-cost" and those less than that as "low-cost." The average constant for a low-cost fund is -0.7 basis points, while the average constant for a high-cost fund is -3.9 basis points. Since the low-cost funds are generally passively managed index-type funds, we find that it isn't surprising for the estimated constant for these funds to be tightly grouped around zero. From the figure, we observe that the large majority of low-cost funds, especially those that charge less than 20 basis points annually, have constants that are within 2 basis points of zero on a monthly basis. On the other hand, we observe a distribution of constants for the high-cost funds that is highly dispersed. Of the high-cost funds, 25% have constants greater than 0, and have a range of -16.7 basis points to +4.7 basis points. All of the TDFs with a constant of negative 10 basis points per month or lower have an expense ratio of 60 basis points per year or higher. There are 35 of them. Almost all of the funds with a constant of more than positive 2 basis points have expense ratios of 30 basis points or higher. We argue that this indicates that the high-cost funds are taking bets on specific securities within the various asset classes. Some sets of bets paid off and some didn't. Since the fitted portfolio represents returns just from taking positions in various asset classes, the constant accounts for how much of the return is due to the fund's individual selections within asset classes, e.g., picking individual stocks or bonds.

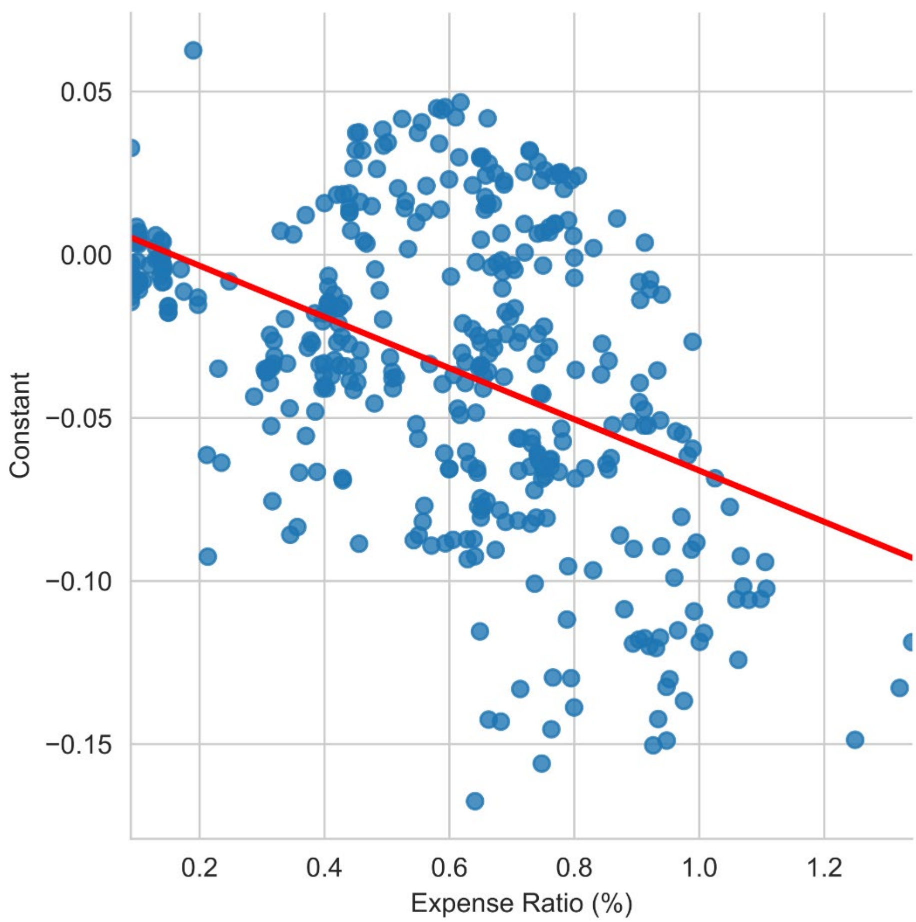

Figure 13: The relationship between net expense ratio of target-date funds and the fitted style analysis constant $\widehat{\alpha}$. The constant scale is in percentage points per month, while the expense ratio scale is in percentage points per year.

A quarter of the high-cost funds are successful in at least meeting the market return, while the other three quarters fall short. Furthermore, the funds that appear to even have a chance of having a positive constant charge less than 80 basis points, with the most successful funds charging around 60 basis points. We think that the negative relationship between constants and expense ratios can be explained by the idea that the various positions taken by actively managed funds will likely aggregate out to be a diversified portfolio, which will on average deliver the market return, scaled by the amount of systematic risk taken on, but with the additional disadvantage of removing high expenses from the fund's return. From the fitted slope of -7.9 basis points in the monthly constant per 1 percentage increase in expenses, this translates to approximately -95 basis points per year. This means that the constant decreases approximately 1 for 1 with an increase in net expense ratio. A 95% confidence interval for the yearly value of -95 basis points contains -1, so we cannot rule out the hypothesis that, on average, once we control for asset allocation, high-cost funds produce gross returns equal to their low-cost counterparts, but net returns that differ by the difference in net expense ratios.

We also run the analogous exercise for the set of balanced funds, and the results are displayed in Figure 14. Balanced funds have a wider range of constants in general than the target-date funds. This is likely due to the class of balanced funds capturing a wider range of funds with different objectives, including varying active strategies, while the target-date funds have a narrower set of objectives. The low-cost funds exhibit a greater clustering around zero than high-cost funds, following the same pattern as the target-date fund data. The regression of the constant on an intercept and expense ratio for the balanced fund data results in $Constant = 0.024 - 0.105 * ExpenseRatio$. Comparing this to the TDF result, the balanced funds have a greater constant by 1.2 basis points, while the slope coefficient is less by 2.5 basis points. However, the error bounds on these estimates do not rule out that the balanced funds and target-date funds have the same intercept and coefficient, as confirmed by a Chow test.

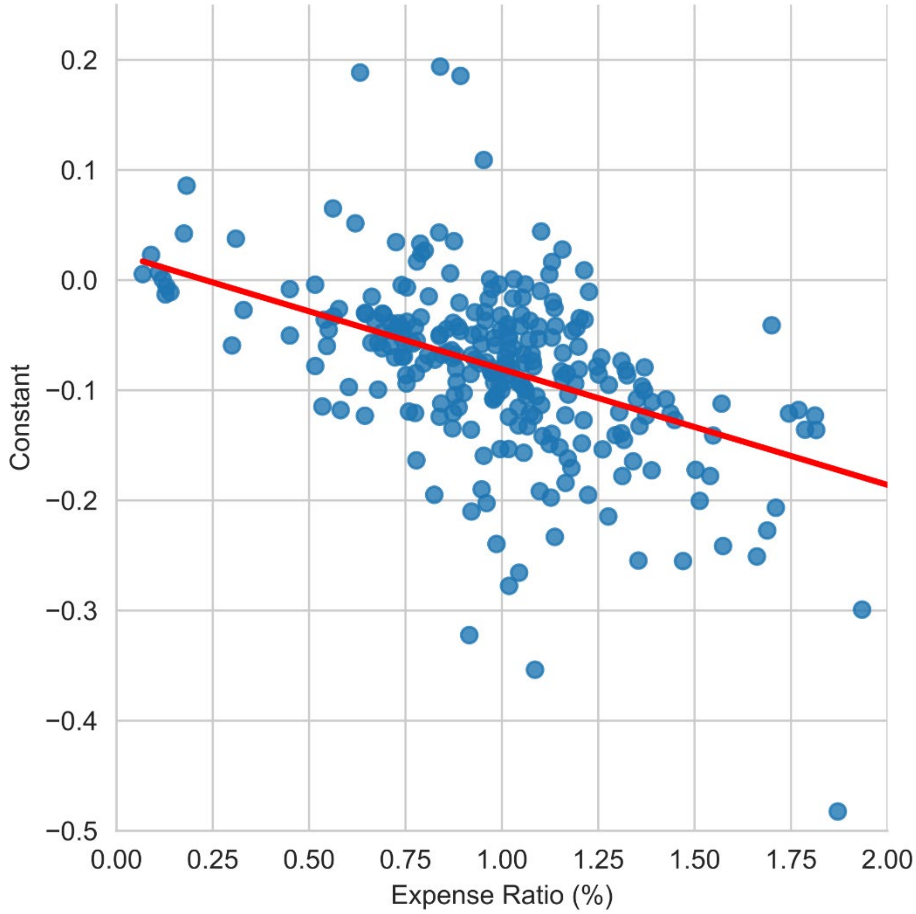

Figure 14: The relationship between the net expense ratio and the fitted style analysis constant $\widehat{\alpha}$ for balanced funds.

8. Return vs. Risk and Longer-Term Performance

In this section, we analyze the performance of target-date funds over the 5-year period beginning in January 2015 going through December 2019. We first look at the return vs. risk tradeoff on a monthly basis through the empirical Sharpe ratio. The empirical Sharpe ratio is constructed by first computing excess monthly returns for all funds, which is the monthly return minus the risk-free return. We used the monthly return on the 3-month Treasury bill as a measure of the risk-free return. We then calculated the empirical Sharpe ratio as the ratio of the mean excess return to the standard deviation of the excess return on a monthly basis over the 5-year period. The results are plotted in Figure 15 against the net expense ratio. The red line is a regression fit, and shows a significant downward trend. For reference, we have included the Sharpe ratios calculated for three ETFs that are market benchmarks: the S&P500 (VOO ticker) in light blue, the total bond market (BND ticker) in green, and the total stock market (VTI ticker) in purple.

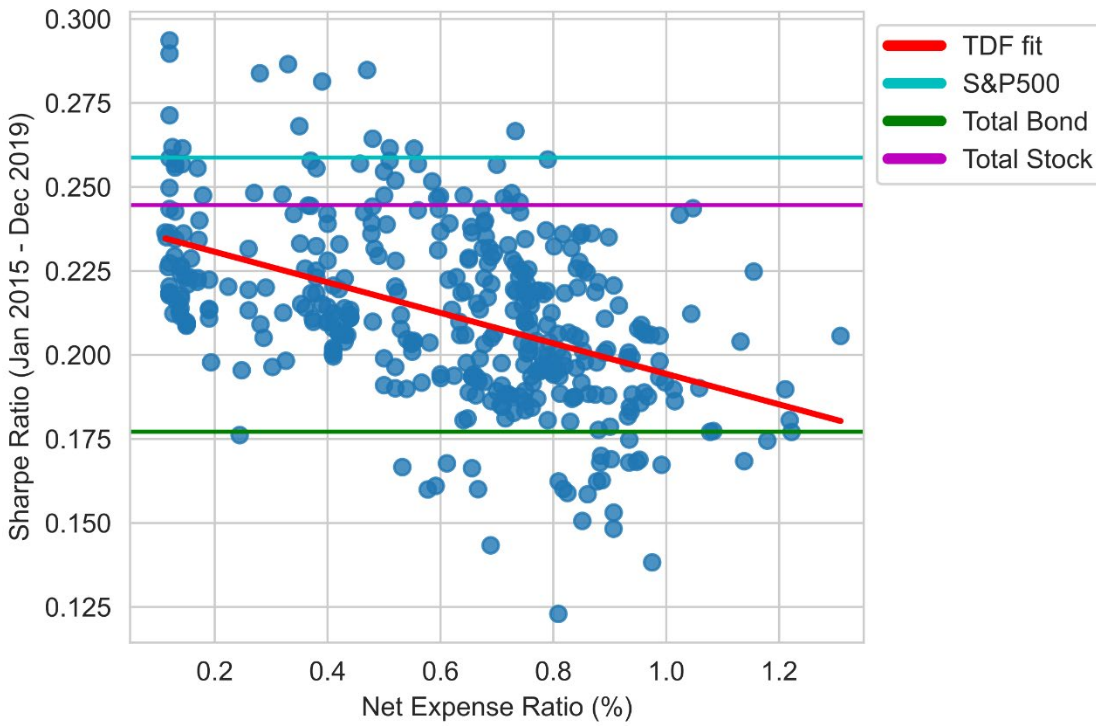

Figure 15: The relationship between the net expense ratio and the Sharpe ratio for target-date funds.

Since target-date funds are mostly a combination of stocks and bonds in terms of the style analysis of previous sections, we expect that the Sharpe ratios for most target-date funds will be somewhere between the total bond and total stock Sharpe ratio values. This is validated in Figure 15, as 82% of the target-date funds in the dataset[3] fall between the green and purple lines. The plot demonstrates once again that it is difficult for expensive funds to overcome their fee disadvantage. As the Sharpe ratio is a measure of excess return per unit of standard deviation, the high-expense funds with the same level of risk as low-expense funds must provide returns that exceed the market return by at least the expense ratio, and it appears that

[3] Funds that did not exist over the entire period Jan 2015 to Dec 2019 were excluded from the Sharpe ratio analysis, just as in the case of the style analysis.

they are unable to do so on average. The slope of the regression line is -0.045 (significant at the 99% level). The interpretation of this is that an increase in net expense ratio by 1%, measured on an annual basis, or about 8 basis points per month, results in an expected loss of Sharpe ratio of 0.045, which is a loss of 4.5 basis points in excess return per 1% of standard deviation. The average standard deviation of monthly return is approximately $\sigma = 2.35\%$. We can calculate the average difference in excess return between two funds that have the average standard deviation but expense ratios that differ by 100 basis points as

$$\frac{1}{\sigma}\left(ExcessRet_{low} - ExcessRet_{high}\right) = SharpeRatio_{low} - SharpeRatio_{high} = 0.045$$

$$ExcessRet_{low} - ExcessRet_{high} = 0.00105.$$

So the difference in excess returns is expected to be about 10.5 basis points, which is not statistically distinguishable from 8 basis points, which is the decrease in expected return due to the expense ratio. This is a result similar to the test from the style analysis; we cannot conclude that the excess gross returns (before fees) of high- and low-cost funds differ systematically.

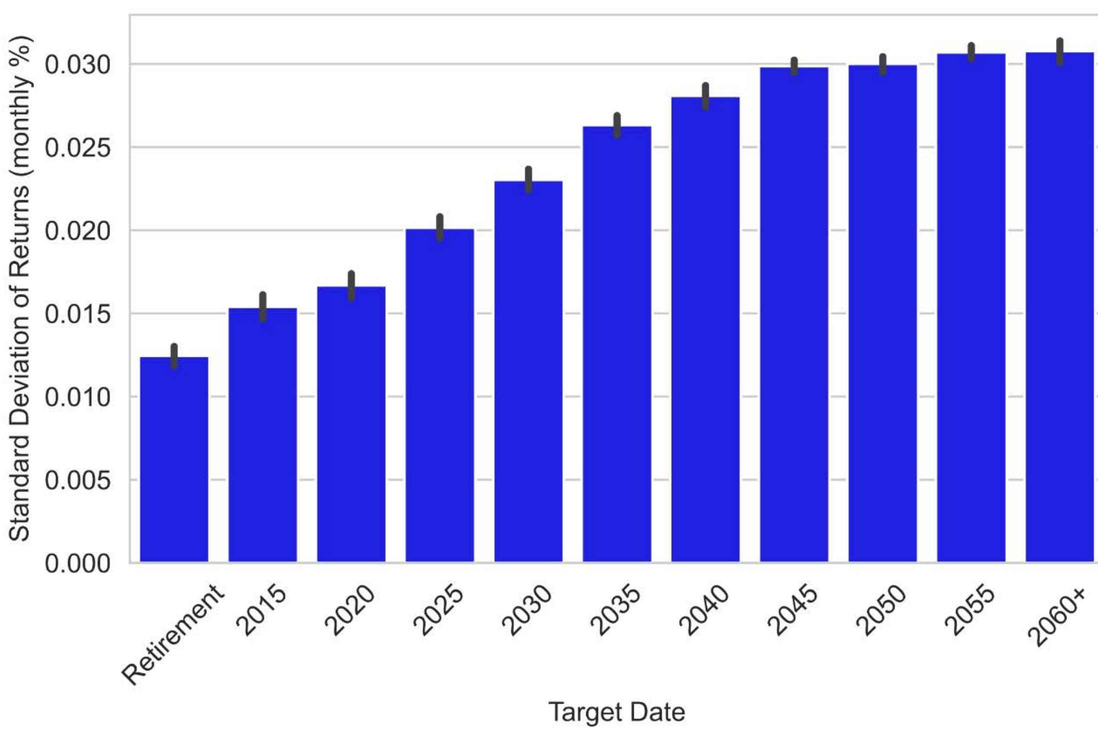

Figure 16: Average of monthly standard deviations within target-date vintages (Jan 2015-Dec 2019). The wicks show the 5^{th} to 95^{th} percentile range.

Figure 16 is another visualization of the glide path present in the target-date fund vintages. The bars roughly follow the same pattern of equity portions in Figure 9. This is perhaps a more direct verification of the goal of target-date funds: that they decrease the riskiness in returns as one approaches retirement. Notably, the standard deviation of funds with vintages 2040-2060+ are not all that different; most of the slope of the glide path occurs when one gets within 15 years of retirement.

We now examine the determinants of future performance for target-date funds. In essence, we are asking the question, "given what I can observe about a target-date fund today, which fund has the best 5-year expected future returns?" In particular, we have data on past returns, expense ratios, assets under management, vintage of fund, and factors obtained from style analysis. We are particularly interested in whether or not past returns have predictive power for future returns, and whether or not the negative effect of expense ratios persists when we control for all of these dimensions. We calculated 5-year annualized returns between January 2010 and December 2014 as the past return ($R10$), and 5-year annualized returns between January 2015 and December 2019 as the future return we are trying to predict ($R15$). The data used for this analysis thus only includes funds that existed over the entire 10-year period 2010-2019.

No. Observations	232
Log-Likelihood	-112.88
Adj. R-squared	0.919

Variable	**Coefficient**	**Standard Error**	**Z statistic**	**P-Value**	**2.5 Percentile**	**97.5 Percentile**
Intercept	5.2017	1.129	4.6074	0	2.989	7.4145
Return 2010-2015	0.0897	0.0329	2.7269	0.0064	0.0252	0.1542
Expense Ratio	-1.0478	0.1229	-8.5251	0	-1.2887	-0.8069
Log AUM	0.0703	0.0198	3.5558	0.0004	0.0315	0.109
Short Bond	-0.0365	0.0123	-2.9591	0.0031	-0.0607	-0.0123
Intermediate Bond	-0.0291	0.0137	-2.1254	0.0336	-0.0559	-0.0023
Long Bond	-0.0094	0.0178	-0.526	0.5989	-0.0443	0.0256
Corporate Bond	-0.0181	0.0117	-1.5421	0.123	-0.0411	0.0049
MBS	-0.0407	0.0119	-3.4094	0.0007	-0.0641	-0.0173

REIT	-0.0409	0.0428	-0.9547	0.3397	-0.1248	0.043
Value Equity	0.0042	0.0147	0.2856	0.7752	-0.0245	0.0329
Growth Equity	0.0402	0.0147	2.7291	0.0064	0.0113	0.0691
Mid Cap Equity	0.0172	0.0169	1.0166	0.3093	-0.016	0.0504
Small Cap Equity	0.0285	0.0148	1.9193	0.0549	-0.0006	0.0576
European Equity	0.0477	0.0186	2.5662	0.0103	0.0113	0.0841
Asian Equity	0.0414	0.0187	2.215	0.0268	0.0048	0.078
Target 2020	-0.1997	0.1057	-1.8896	0.0588	-0.4068	0.0074
Target 2025	-0.1909	0.1177	-1.622	0.1048	-0.4215	0.0398
Target 2030	-0.2397	0.149	-1.6081	0.1078	-0.5318	0.0524
Target 2035	-0.3629	0.1838	-1.9741	0.0484	-0.7232	-0.0026
Target 2040	-0.3582	0.2116	-1.6927	0.0905	-0.773	0.0566
Target 2045	-0.409	0.2447	-1.6713	0.0947	-0.8887	0.0707
Target 2050	-0.345	0.2466	-1.3987	0.1619	-0.8283	0.1384
Target 2055	-0.1338	0.2662	-0.5027	0.6152	-0.6556	0.388
Target Retirement	0.0144	0.0989	0.1461	0.8838	-0.1793	0.2082

Table 3: Regression of annualized Jan 2015-Dec 2019 return on predictors with robust standard errors.

We specify the regression

$$Ret15_i = c + \alpha Ret10_i + \beta ExpRatio_i + \gamma\, log(AUM_i) + \boldsymbol{\rho F_i} + \boldsymbol{\tau V_i} + \epsilon_i$$

Where $\boldsymbol{F_i}$ is a vector of factors determined from the style analysis fit over 2010-2014 and $\boldsymbol{V_i}$ is a vector of dummies indicating the vintage of the fund. All of the explanatory variables are information available to an investor at the beginning of 2015, so the expense ratios and AUM given are as of the beginning of 2015. We ran the regression with weights determined by the AUM, and computed heteroscedasticity-robust standard errors. The results are presented in Table 3. We find that the previous 5-year return has a statistically significant coefficient of 0.0897. This can be interpreted as the increase in expected return from having a 1% higher past return, so a 1% higher return per year over 2010-2014 translates to an increase in expected return of about 9 basis points (annualized). This indicates large regression to the mean for past good performers, though there is still a small effect. The expense ratio coefficient is -1.0478, which a z-test confirms is indistinguishable from -1. This is one more test that gross returns (before expenses) from high- and low-cost funds are on average the same, though as in the

Sharpe ratio analysis, the estimate assigns a higher expected gross return to low-cost funds. The log AUM also has a small but statistically significant positive effect on expected returns, a possible indicator that larger funds may have slightly better information or more talented management. The factor coefficients have the signs one would expect, negative for bonds and alternatives, and positive for equities, though they are not statistically significant. They are interpreted as the increase in expected return from a 1% increase in portfolio weights to the factor. Finally, the target date dummies are almost all insignificant, since once we control for fund factors, it is reasonable that all target-date funds are similar.

We do one more exercise to demonstrate the differences between high- and low-cost funds. Given that past returns are small but positive predictors of future returns, how much better would a high-cost fund have to have performed in the past to justify its investment over a low-cost fund? The AUM-weighted average expense ratio is 13.3 basis points for the low-cost funds, and 63.7 basis points for the high-cost funds. Keeping all else equal, for a low-cost fund with 13.3 basis point expense ratio and a high-cost fund with 63.7 basis point expense ratio to have the same predicted 5-year return, the required difference in the previous 5-year return is $Ret10_{high} - Ret10_{low} = \frac{1.0478(0.637-0.133)}{0.0897} = 5.89\%$. For a 50-basis point gap in expense ratios, then, the expensive fund needs to perform almost 6% higher annually for the past 5 years to close the gap.

9. One Size Fits All or One Size Fits Almost Nobody

The asset allocation of target date funds in defined contribution plans takes one factor into account: the employee's age. While age is an important determinant of risk aversion, it is not the only one by any means. At any given age, people would exhibit a wide range of risk tolerance. Some of this is simply because people are different in terms of their preferences and some of it is that people face different situations. Their circumstances would differ in important ways such as marital status, whether they have children, spouses' employment, whether they anticipate supporting their parents in old age or receiving an inheritance from them, their financial assets outside of employer-sponsored retirement accounts, the risk characteristics of their profession (their human capital), their education, their homeownership status, and their financial liabilities such as student loans.   None of these factors are taken into account in plans which default workers into target date funds based solely on age.

The recent paper by Mitchell and Utkus (2020) finds that participants in plans where low-cost TDFs are introduced as the default option improve their asset allocation in a number of important ways. They increase their allocation to equities and bonds and decrease the allocation to company stock and cash, for instance. While their evidence indicates improved investment exposures, it does not make the case that target date funds are the optimal default investment for retirement saving.

The case for the kind of glide path offered by TDFs is best made by Bodie, Merton and Samuelson (1992). They argue that a proper life cycle model would account for human capital as well as financial capital. Human capital, the present value of future anticipated labor market earnings, naturally declines as one's career progresses and as retirement approaches. When young, human capital is likely the predominant form of wealth. If one's human capital is low risk or has riskiness similar to high-grade bonds, then it may make sense to have a very large fraction of financial wealth in risky categories such as equity.   As human capital becomes less important with age, it would then make sense to increase bond exposure in financial assets such as 401(k) accounts. While this logic may support an asset allocation glide path, it certainly does not support the one size fits all nature of target date funds. Some forms of human capital may have equity-like risks (for instance, real estate agents or tech workers, whose compensation features bonuses and equity grants) whereas other careers may generate low risk human capital (e.g. tenured public school teachers). Bodie, Merton and Samuelson's work suggests that the risk characteristics of one's human capital are an important determinant of optimal asset allocation in financial investments.

No. Observations	465
Log-Likelihood	-609.87
Adj. R-Squared	0.72

Variable	Coefficient	Standard Error	Z statistic	P Value	2.5 Percentile	97.5 Percentile
Intercept	3.6724	1.4204	2.5856	0.0097	0.8886	6.4563
TDF Indicator	0.1327	0.1045	1.2699	0.2041	-0.0721	0.3374
Return 2010-2015	0.0563	0.1176	0.4788	0.6321	-0.1741	0.2867
Expense Ratio	-1.0204	0.1463	-6.9757	0	-1.307	-0.7337
Log AUM	0.0053	0.0163	0.3251	0.7451	-0.0267	0.0373
Short Bond	-0.0132	0.0137	-0.9634	0.3353	-0.04	0.0136
Intermediate Bond	0.0056	0.0171	0.327	0.7437	-0.028	0.0392
Long Bond	0.0217	0.0253	0.8566	0.3917	-0.028	0.0714
Corporate Bond	-0.0026	0.0144	-0.1791	0.8578	-0.0307	0.0256
MBS	-0.0164	0.0163	-1.0033	0.3157	-0.0483	0.0156
REIT	0.1334	0.0652	2.0443	0.0409	0.0055	0.2612
Value Equity	0.0445	0.0174	2.5666	0.0103	0.0105	0.0786
Growth Equity	0.0913	0.0166	5.5013	0	0.0588	0.1238
Mid Cap Equity	0.0316	0.0228	1.3882	0.1651	-0.013	0.0763
Small Cap Equity	0.0207	0.0157	1.3228	0.1859	-0.01	0.0514
European Equity	0.064	0.0236	2.7078	0.0068	0.0177	0.1104
Asian Equity	0.0037	0.0201	0.1845	0.8536	-0.0357	0.0432

Table 4: Regression of Annualized 2015-2019 return on pooled target-date fund and balanced fund data controlling for fixed effects of target-date funds (with robust standard errors).

There is an alternative to target date funds as the default option in defined contribution retirement plans. That is to use their close relative, traditional balanced funds, which are also known as target risk funds. It is then worth asking if balanced funds and target-date funds with similar investment profiles differ systematically in terms of returns. Table 4 reports the results of a regression specification similar to that of Table 3 where both target-date and balanced fund data are pooled and an indicator is included for target-date funds. The coefficient on this indicator is positive but small and statistically insignificant, which serves as a test of systematic differences between target-date funds and balanced funds with similar investment profiles, and we conclude that there is not sufficient evidence to say they are different.

An alternative to a target date fund as the default option in an automatic enrollment retirement plan could be the moderate risk balanced fund with equity exposures ranging from 50 to 70 percent. But, all five target risk categories (very conservative, conservative, moderate, aggressive and very aggressive) would be available in the defined contribution plan. Employees could be strongly encouraged to take a risk tolerance questionnaire every five years to determine which of the target risk funds best matches their circumstances. By using two of the target risk offerings, employees could precisely choose their desired level of equity exposure. If this became the norm, it is almost certain that advice services, both online and in person, would become available. The advantage of this approach is that individual employees could take into account their personal circumstances in choosing their retirement portfolio and the glide path of their asset allocation over their working life. Potentially, this would replace a one size fits all default offering with a custom-fit portfolio. We are not advocating this switch in default options, simply suggesting that it might be worth further research and consideration.

10. Conclusion

While finance economists may have found little surprising in our in-depth analysis of target date funds, we think some participants and 401(k) and 403(b) committee members at employers may find some of the findings useful. Perhaps the biggest takeaway is that even the near term target date funds still have considerable stock market exposure. For instance, we find that 2025 target date funds lost between 20 and 25 percent of their value in five weeks in February and March, 2020. Our guess is that these losses were larger than anticipated by either participants or 401(k) and 403(b) committee members.

We find that the average higher-cost actively managed target date funds failed to perform as well as the cheaper indexed competition in the 2015-2019 period. Some of the actively funds did very well in relative terms, but most did not. We found that past performance is only weakly predictive of future performance. The implication is that even an active fund with a superior record has an expected future return below the passive alternative TDFs. These results may be specific to the 2010-2020 period examined, but only time will tell. We know that this was a period where active funds, in general, had difficulty matching passive funds. It is also possible that some actively managed TDFs can systematically outperform the market and the competition. The data indicate that such performance would be an outlier.

The last section of the paper discussed the one size fits all nature of target date funds and suggests that a traditional family of balanced funds (AKA target risk funds) offer an alternate default approach. One thing should be emphasized - even though there are a lot of flavors of target date funds, an employee participating in a defined contribution plan typically only is offered one such family of funds. The 2025 vintage might feature a particular level of equity exposure between 50 and 65 percent. It is not hard to imagine circumstances for an employee where the optimal equity exposure at age 60 would be 20 percent (or 80 percent). This illustrates the drawback of a one size fits all offering such as a target date fund.

Target date funds have clearly passed the market test in that they have attracted enormous investments over the past twenty years. Their size alone merits monitoring their performance. Markets provided a natural stress test in February and March of 2020. Even the short duration TDFs suffered large valuation losses over a five week period. Fortunately, markets completely recovered in the subsequent three months, so few participants suffered substantial losses. But, the episode did reveal that TDFs expose even those near retirement to substantial market risks.

References

Bodie, Zvi, Robert Merton, and William Samuelson, “Labor Supply Flexibility and Portfolio Choice in a Lifecycle Model,” *Journal of Economic Dynamics and Control*, Vol 16, pp 427-449, 1991.

Elton, Edwin J, Martin J. Gruber, Andre de Souza, and Christopher R. Blake, “Target Date Funds: Characteristics and Performance” (March 23, 2015). Available at SSRN: https://ssrn.com/abstract=2547558 or http://dx.doi.org/10.2139/ssrn.2547558

General Accountability Office (GAO), “Defined Contribution Plans: Key Information on Target Date Funds As Default Investments Should Be Provided to Plan Sponsors and Participants,” GAO-11-118, January 2011. https://www.gao.gov/new.items/d11118.pdf

Gomes, Francisco J., Laurence J. Kotlikoff and Luis M. Vicelra, “Optimal Life-Cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life-Cycle Funds,” *American Economic Review*, May 2008.

ICI, Release*: Quarterly Retirement Market Data, Fourth Quarter 2019*. Washington, DC: Investment Company Institute, 2020. https://www.ici.org/research/stats/retirement

ICI, *2020 Investment Company Fact Book*, Washington, DC: Investment Company Institute, 2020. https://www.ici.org/pdf/2020_factbook.pdf

Merton, Robert, “Lifetime Portfolio Selection Under Uncertainty: The Continuous Time Case,” *Review of Economics and Statistics* Vol. 51: pp. 247-57, 1969.

Mitchell, Olivia S. and Stephen Utkus, “Target-Date Funds in 401(k) Retirement Plans,” NBER Working paper no. 17911, March 2012.

Mitchell, Olivia S. and Stephen Utkus, “Target Date Funds and Portfolio Choice in 401(k) Plans,” NBER working paper No. 26684. January, 2020

Poterba, James, Joshua Rauh, Steven Venti, and David Wise. “Utility Evaluation of Risk in Retirement Savings Accounts,” In David Wise (Ed). *Analyses in the Economics of Aging*, Chicago: University of Chicago Press, 2005.

Poterba, James, Joshua Rauh, Steven Venti, and David Wise, “Lifecycle Asset Allocation Strategies and the Distribution of 401(K) Retirement Wealth.” In David Wise (Ed). *Developments in the Economics of Aging*, Chicago: University of Chicago Press, pp. 333-79, 2009.

Samuelson, Paul A., “Risk and Uncertainty: The Fallacy of The Law of Large Numbers,” *Scientia* 98: 208-113, 1963.

Samuelson, Paul A., “Lifetime Portfolio Selection by Dynamic Stochastic Programming,” *The Review of Economics and Statistics*, Vol. 51, No. 3, pp. 239-46, August 1969.

Sharpe, William F., “Asset Allocation: Management Style and Performance Measurement: An Asset Class Factor Model Can Help Make Order out of Chaos*.” Journal of Portfolio Management*, pp. 7-19, Winter 1992.

Shiller, Robert, “Lifecycle Portfolios As Government Policy,” *Economists Voice*, Vol. 2, pp. 1-9, 2005.