Buy the Dip

2 authors:

Thomas Shohfi
Rensselaer Polytechnic Institute

Majeed Simaan
Stevens Institute of Technology

Buy the Dip*

Thomas Shohfi†
shohft@rpi.edu

Majeed Simaan‡
msimaan@stevens.edu

April 27, 2021

Abstract

The increasing role of social media in financial markets has encouraged retail traders to "buy the dip" (BTD). We present a simple paradigm describing this strategy in terms of dip size and purchase smoothing. The empirical investigation considers different specifications and testing periods. While BTD does not necessarily maximize investors' terminal wealth and is sensitive to market conditions at the beginning year of investment, it does provide a heuristic approach to improve risk-adjusted performance over a passive investment policy. Overall, BTD provides a simple, intuitive approach in dealing with portfolio selection over time.

Keywords: *Asset Allocation, Bear/Bull Market, Market Timing, Social Media*

JEL Codes: *G10, G11, G17, G41*

***Disclaimer**: The authors are providing information, but make no direct fiduciary recommendations on its implementation. This paper to be used by an investor by their own choice and merits. One should consult a financial advisor if needed.

†Lally School of Management and Technology, Rensselaer Polytechnic Institute, 110 8th St, Troy, NY 12180, USA.

‡School of Business, Stevens Institute of Technology, 1 Castle Point Terrace, Babbio Center, Hoboken, NJ 07030, USA.

1 Introduction

The term “buy the dip” (henceforth BTD) has become popular on social media platforms, especially among millennial retail traders. For example, Google search activity for the term “buy the dip” surged over 1500% during the March 2020 coronavirus-driven market selloff. The median number of tweets containing the term is about 581 per day, according to a total of 5,319 tweets between March 25, 2021, and April 2, 2021. For instance, Figure 1 illustrates the intraday prices of the SPY ETF along with the number of tweets per 15 minutes increments.[1] The number of tweets denotes how many cases covering the $SPY “cashtag” (Cookson and Niessner (2020)) occur and, simultaneously, include a dip-related term. The figure illustrates the increase of dip-related tweets along with intraday SPY price movements.

Figure 1: **Buy the Dip Tweets and SPY Intraday Prices**
The following plot illustrates the number of tweets containing the $SPY “cashtag” and mentioning a dip-related term (blue line). For each day, this number corresponds to the total cumulative tweets. For instance, on April 6, 2021, the series resets and considers the cumulative number over that day rather than the previous day. The orange line denotes intraday prices of SPY. In either case, the figure illustrates results using 15 minutes increments. Twitter data is collected using the Twitter open API and downloaded using the `twitteR` R library by Gentry (2015). The intraday stock data is collected from Alpha Vintage according to Dancho and Vaughan (2020).

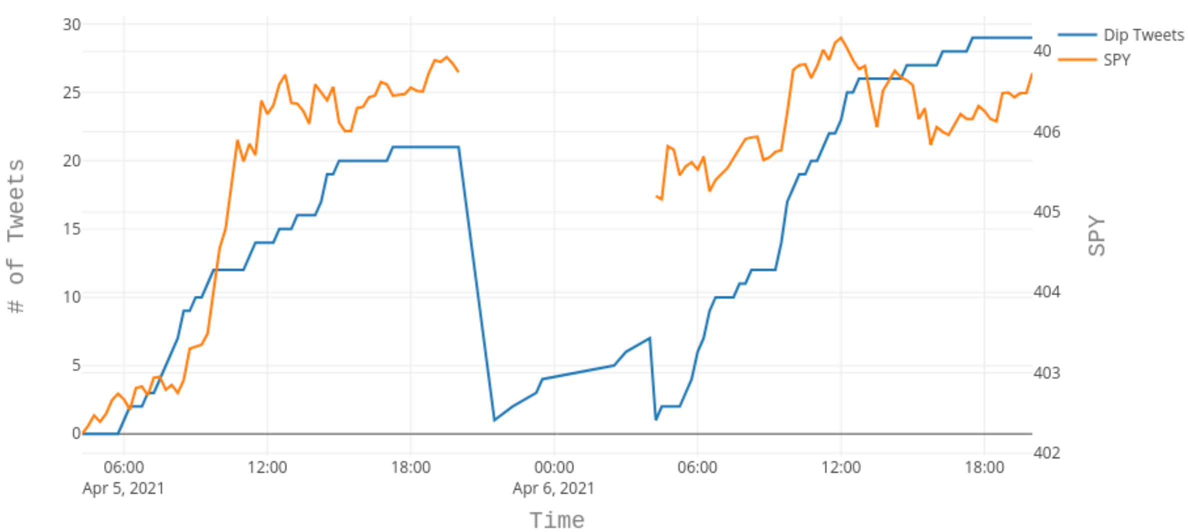

Figure 1 provides a simple anecdote to tie the appeal of BTD on social media and the price movement of SPY. In terms of asset allocation, a BTD policy resembles a market-timing approach in which the investor waits for a specific contingency to take place before allocating more wealth to the risky asset. However, such timing also comes with an opportunity cost. Investors could forego investment opportunities with relatively high-risk premiums while holding cash for more extended periods. At the same time, selloffs could also reflect the persistence of bad news in the market. As a result, accumulating cash during bull periods to be invested following significant market drops may not reflect the most optimal investment policy. Nevertheless, behavioral studies recognize that bad news has a stronger effect than good news (Baumeister et al. (2001)). Hence, if market participants

[1] This illustration is conducted over two days, covering April 5th and 6th of 2021.

overreact to bad news, BTD could potentially lead to a correction in the market and constitute an optimal investment.

This research investigates the appeal of BTD in terms of optimal strategic asset allocation. According to Sharpe (2010), a strategic asset allocation denotes a process with well-defined steps. Following these steps results in a specific allocation of capital to the underlying assets, i.e., the policy portfolio. In this regard, we follow a set of rules ex-ante and investigate whether BTD corresponds to an optimal policy portfolio ex-post. In particular, such policy determines the allocation of wealth to the risky asset for a given period, under certain market conditions, and according to specific budget constraints. We evaluate the performance of a given BTD policy versus a passive one as the primary benchmark. By passive strategy, we denote the case in which the investor allocates wealth to the risky asset as soon as funds become available. Overall, our analysis concerns the investment in the market portfolio, proxied by the SPY ETF.

We consider two main specifications to evaluate BTD policies. In the first case, we consider a lump sum case in which the investor allocates an endowment over time without any future cash flows. In the second one, we consider the case in which the investor receives fixed cash flows over time. In either case, the investor allocates a specific unit of capital only when the market drops by a certain threshold. Our paper makes a couple of important contributions. First, whether BTD policy enhances the terminal wealth of the investor over time is subjected to several factors: (i) Market conditions at the beginning of the period, i.e., bull versus bear market; (ii) The threshold that the investor considers constituting the dip; and (iii) the measurement of the dip. The first factor is likely to be out of the investor's control and highly dependent on his/her birthdate (Dimson et al. (2021)). However, there is much potential improvement of risk-adjusted performance when investors use maximum drawdown (MDD) to measure potential dip.[2]

Second, while mean-variance optimization is designed to improve investors' risk-adjusted returns, it is also subject to estimation risk that hinders its ex-post optimality (see, e.g., Michaud (1989); Best and Grauer (1991); Kim et al. (2015)). Given such challenges in a dynamic setting (DeMiguel et al. (2015)), our findings show that BTD is a heuristic approach to improve risk-adjusted returns over time.[3] This result is consistently evident in the case of monthly cash inflows. While investors may not maximize their terminal wealth using BTD compared to a passive policy, they can attain a higher risk-adjusted reward with lower downside risk by buying the dip.

Under the "market efficient hypothesis" (Fama et al. (1969)), returns on the market follow a random walk and, hence, are unpredictable. Most finance academic researchers believe that the markets follow weak-form efficiency (Doran et al. (2010)). Nonetheless, the behavioral finance literature suggests that returns exhibit a form of predictability that can be attributed to predictable investor behavior such as overreaction and overconfidence (De Bondt and Thaler (1985)).[4] A BTD strategy is inherently contrarian, which is more common than a momentum strategy among retail

[2]Risk-adjusted performance is measured with respect to the Sortino ratio (Sortino and Van Der Meer (1991)).

[3]For instance, Wang and Taylor (2018) find weak evidence in favor of a dynamic portfolio policy over a static portfolio one.

[4]For a more recent discussion see, e.g., Golez and Koudijs (2018).

investors. These long-term contrarian retail investors have asymmetric market-timing preferences as they are more likely to "buy the dips" instead of "sell the rips" (Goetzmann and Massa (2002)). While our paper considers the portfolio strategy from the perspective of retail investors rather than fund managers, the argument of whether investors should pursue a BTD policy is similar to the debate whether active managers are capable of outperforming a passive fund/strategy. Overall, our findings indicate that BTD is a heuristic tool that retail investors could benefit from to attain higher risk-adjusted returns over time.

Our paper has important implications for robo-advisers and their impact on retail investors (Loos et al. (2020)). BTD is subjected to less estimation error and, hence, model risk. In this regard, BTD provides better transparency into decision-making than predictive models that rely on historical data (Clark et al. (2020)). Additionally, a BTD policy provides a simple yet intuitive approach that helps investors allocate their wealth over time with better explainability. At the same time, investors who are less risk averse could be better off investing their cash flows as soon as they become available to maximize their terminal wealth. Overall, our results indicate that waiting for the dip could enhance investors' risk-adjusted returns, however, at the cost of lower terminal wealth.

The manuscript proceeds as follows. Section 2 discusses the main investigation behind our analysis. In Section 3, we discuss the main results and findings of the paper. Finally, Section 4 concludes.

2 Investigation

2.1 Data

The data set corresponds to two time series. The first is the daily return on SPY, the original S&P 500 exchange-traded fund (ETF). The series dates between Jan 1994 and Dec 2020. The other series corresponds to the monthly Consumer Price Index (CPI) collected from the Federal Reserve Economic Data (FRED). We use the latter to measure inflation, which is computed as the monthly percent change in CPI. Based on the monthly time series, we adjust stock returns accordingly. Let R_{t+1} denote the return on the portfolio between month t and $t+1$, whereas i_{t+1} denotes the inflation rate for the same period. The inflation-adjusted monthly return is

$$\tilde{R}_{t+1} = \frac{1 + R_{t+1}}{1 + i_{t+1}} - 1 \tag{2.1}$$

In Figure 2, we illustrate the cumulative return on SPY over the whole period. If one bought SPY at the end of 1993, one million dollars would have become \$13.24 million as Panel (a) illustrates. In Panel (b), we provide a similar perspective, however, by taking into considerations the inflation-adjusted return with respect to Equation (2.1). Adjusted to inflation, Panel (b) indicates that a passive investment in SPY would have grown a \$1 million investment to \$7.4 million.

Figure 2: **SPY Cumulative Return**
This figure illustrates the cumulative return of the SPY ETF over time. The full sample period dates between Jan 1994 and Dec 2020. Panel (a) denotes raw-returns over the whole period. Panel (b) considers inflation-adjusted returns with respect to Equation (2.1). Grey bars denote recession periods identified with respect to the National Bureau of Economic Research (NBER).

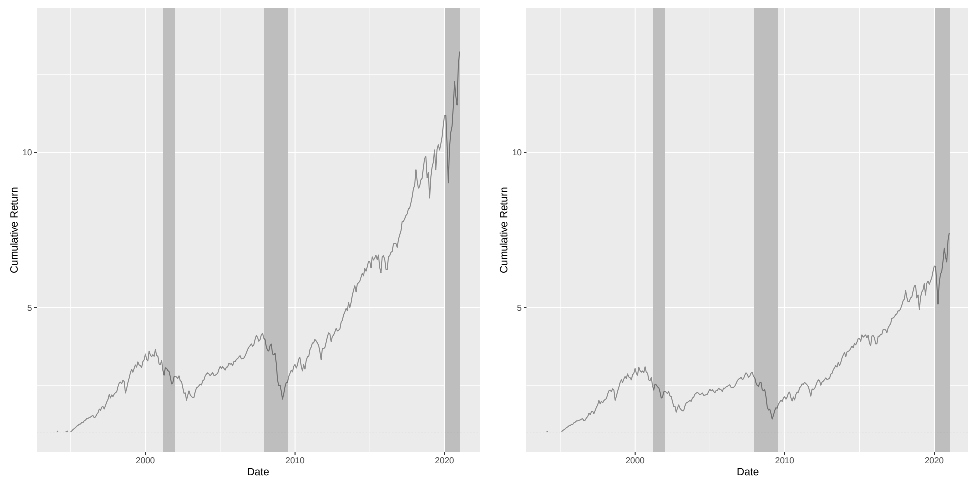

(a) Raw Returns (b) Inflation-adjusted

2.2 Portfolio Selection

2.2.1 Challenges

The result from Figure 2 denotes the return on a lump sum investment strategy. That is the case for investors who allocate their whole wealth to the risky asset as soon as trading becomes possible. Should investors allocate this wealth differently over time? Let us consider this problem from a portfolio point of view. In particular, let θ_t denote the proportion of wealth allocated to the risky asset at time t. The total wealth of the investor is w_t, such that $w_0 = 1$ denotes the lump sum endowed at time $t = 0$. Given the notation of inflation rate and inflation-adjusted return, the wealth over the next period is given by

$$w_{t+1} = \theta_t w_t (1 + \tilde{R}_{t+1}) + (1 - \theta_{t+1}) w_t \frac{1}{1 + i_{t+1}} \tag{2.2}$$

The above can be rewritten as

$$w_{t+1} = \frac{w_t}{1 + i_{t+1}} \left[1 + \theta_t R_{t+1}\right] \tag{2.3}$$

Suppose that the investor has a mean-variance preference and the allocation decision is concerned with one period. Also, assume that inflation over this period is constant. Hence, the optimal choice of θ_t can be written as

$$\theta_t^* = \frac{1}{A} \frac{\mu_{t+1|t}}{\sigma^2_{t+1|t}} \tag{2.4}$$

where $\mu_{t+1|t}$ and $\sigma_{t+1|t}$ denote the conditional expected return and volatility of the risky asset over the next period, whereas A denotes the risk aversion level of the investor.

While Equation (2.4) has a normative appeal in terms of taking into account risk-reward trade-off, it has two major drawbacks. The first one is related to the fact that the solution is myopic in the sense that it does not take into consideration future hedging demand. Such future demand is determined by news about future returns on invested wealth (Merton (1973)). A non-myopic consistent mean-variance solution can be attained with respect to Basak and Chabakauri (2010). Nonetheless, such solution is subject to estimation risk, which brings us to the second drawback. In practice investors need to estimate the corresponding parameters $\mu_{t+1|t}$ and $\sigma_{t+1|t}$ for the choice to become feasible.[5] It is well-documented that estimation error in mean-variance portfolio selection is severe and leads to poor out-of-sample performance.[6] For example, DeMiguel et al. (2009) highlight the appeal of a naive passive strategy in terms of mean-variance out-of-sample efficiency. While such a strategy does not have the same normative appeal as the mean-variance portfolio, it provides a better out-of-sample performance and potentially leads to higher out-of-sample utility.

2.2.2 Investment Policy

Let $g(\theta; p)$ denote the allocation decisions of policy p over time. By policy, we refer to the set of optimal decisions over time, i.e. $\{\theta_t(p)\}_{t=0}^{T-1}$. For instance, the passive policy that invests full wealth at point $t = 0$ is denoted by $p = N$, such that $\theta_t(N) = 1$ for $t = 0, 1, ..., T-1$. In addition to the naive policy, we consider a number of allocation decisions that represent the BTD idea, which will be discussed shortly. For different policies, we consider a couple of performance criteria:

1. **Terminal Wealth**: In terms of utility, terminal wealth denotes a higher level of future consumption. This is highly relevant for investor's seeking to maximize their retirement savings.

2. **Risk-adjusted Return**: To evaluate the mean-variance efficiency of the portfolio over time, we consider the Sortino ratio (Sortino and Van Der Meer (1991)) of the inflation-adjusted return on wealth over the time period. The measure provides a risk-reward assessment of the underlying policy. Risk is measured by the semi-deviation, i.e. volatility of wealth when it decreases. Reward, on the other hand, is measured by mean return on wealth.

Buying the Dip: Rather than investing the full amount at once, suppose that the investor is willing to allocate her wealth into k sums that will be invested over time. The investor starts with $1/k$ of her wealth in stocks and $1 - 1/k$ in cash. If SPY's monthly return is higher than inflation then the investor faces a positive opportunity cost by holding cash. If SPY's return is negative enough, then the investor could be better off holding cash than stock. Nonetheless, the investor

[5] DeMiguel et al. (2015) propose a shrinkage approach to deal with estimation risk for multiperiod mean-variance portfolio selection, while taking into consideration transaction costs.

[6] See, e.g., Michaud (1989); Best and Grauer (1991); DeMiguel et al. (2009) and references therein.

adds a fraction of her cash to stock each time SPY drops more than a predetermined return denoted by τ.

To better understand the logic behind the above strategy, consider the case for $k = 2$ and $\tau = 0$. In this case, the investor starts with $1M with half in SPY and the rest in cash. She only adds the other half when SPY becomes negative for the first time. Nonetheless, what are the "optimal" values for k and τ? How large a dip should τ be before buying? Perhaps investors have different levels of market experience that shape prior beliefs with respect to what size dip represents an optimal opportunity to buy stocks? (Nicolosi et al. (2009)) This is the idea behind BTD. From a dynamic programming point of view, the choice of k and τ denote a single investment policy. In the analysis below we do not consider how the investor chooses this policy ex-ante. However, we evaluate different policies to investigate the appeal of BTD as an investment policy versus a naive one.

2.3 Implementation

In our empirical analysis, we consider three problems. The first one corresponds to the case in which the investor is endowed with a lump sum that will be invested over time. In the second one we consider the problem of fixed cash flows that the investor receives at the beginning of each month. The allocation of these cash flows to the risky asset are determined by the parameter τ alone. The third problem is a combination between the two. The investor starts with 12-months savings. As time passes, she receives fixed monthly inflows as before. Different from the second approach, she allocates her total cash savings to the risky asset depending on the maximum drawdown (MDD) of SPY. In all cases, cash is adjusted for inflation to take into account the associated opportunity cost. We discuss strategy each below.

Lump Sum

We consider 60 investment policies determined by different combinations between k and τ. We set $k = \{1, 2, ..., 10\}$ and $\tau = \{-0.05, -0.04, ..., -0.01, 0.00\}$. The investigation is conducted on a monthly basis. For instance, for $k = 10$ and $\tau = -0.01$. The investor allocates an additional fraction (one tenth) of her cash to the risky asset when the monthly return on SPY drops below -1%.

Monthly Inflows

Most investors are unlikely to be endowed with a large lump at a single initial point of time and are more likely to invest a percentage of periodic cash flows (i.e., salary payments). In this case, the investor enters the trading period with a single unit of cash. At the beginning of each month, she receives an additional unit of cash.[7] Over the full sample period, the investor receives 324 such inflows. The passive policy allocates cash to the risky asset as soon as funds become available. In

[7] All cash inflows are inflation-adjusted to December 2020 constant dollar values.

other words, she allocates a single unit of cash to the risky asset at the beginning of each month. On the other hand, BTD dictates that she invests a single cash unit to SPY if its daily return drops below τ. This is subject to the constraint that her outstanding funds are greater than zero. Our implementation sets $\tau = \{-0.020, -0.015, -0.010, -0.005, 0.000\}$.

Maximum Drawdown

The investor enters the trading period with 12 units of cash, emulating a one-year savings. At the beginning of each month, she receives an additional unit of cash. Over the full sample period, the investor receives 312 such inflows. Different from the previous strategy, the investors allocates total savings into the risky asset if the MDD of the risky asset is less than τ. By MDD, we use m-months rolling window returns to determine the value. We set m to be either 3, 6, or 12. For the threshold, we set $\tau = -20\%, -10\%, -5\%$ corresponding to the return levels associated with bear markets, corrections, and pullbacks, respectively (Barron's (2019)). If the m-months MDD is less than τ, then the investor allocates her funds to the risky asset. When funds become available in the following month, she allocates the new unit to the risky asset as long as the MDD is less than τ. During periods in which the MDD is not below τ, the investors accumulates cash in anticipation of the next event where the MDD becomes less than τ.

Testing Period

We consider different testing periods. In all cases, we set the terminal date as Dec 2020. The main difference is the starting period. We consider four different starting periods. The first two are associated with a bull market beginning. Those are years beginning Jan 1994 and Jan 2010. The other two periods are Jan 2000 and Jan 2008 that correspond to an initial period of bear markets. In all cases, we compare the investment policies with respect to the passive one.

3 Results and Discussion

The discussion below is devoted to two parts. The first one covers the lump sum problem, whereas the second one corresponds to the fixed cash flow problem.

3.1 Lump Sum

To gain an initial perspective on BTD policy, we consider a special case of $k = 20$ and $\tau = \{-0.06, -0.05\}$ in Figure 3. We observe that most purchases of the risky asset occur during times of distress. For instance, going through the global financial crisis, the investor enters with a 0.2 fraction of her cash when $\tau = -0.05$. This amount eventually is exhausted by the end of 2009. On the other hand, if the investment policy is dictated by $\tau = -0.06$, we note that it is only prior to 2020, when the investor allocates the whole sum to the risky asset. Overall, Figure 3 demonstrates the sensitivity of BTD to selection of dip size τ and frequency k.

Figure 3: **Cash Holdings over Time**
This figure denotes the cash holdings of two investment policies. The investment policies relate to the lump sum problem in which the investor is endowed with a single unit of cash. Panels (a) and (b) illustrate the cash holdings over time by following BTD policy for $\tau = -0.05$ and $\tau = -0.06$, respectively. Grey bars denote recession periods identified with respect to the National Bureau of Economic Research (NBER).

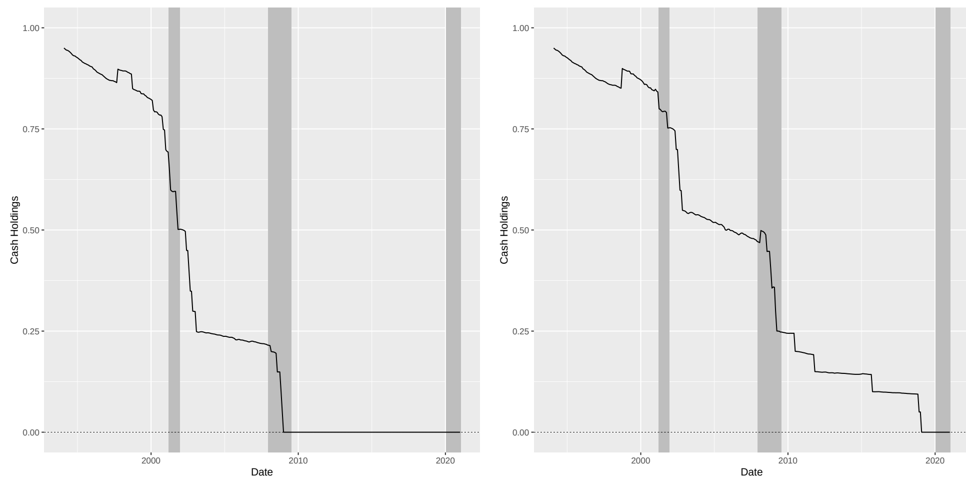

(a) $\tau = -0.05$ (b) $\tau = -0.06$

In Panel (a) from Table 1, we report terminal wealth for different policies discussed in Section 2.3. A couple of comments are in order. First, the case of $k = 1$ denotes the passive investment (i.e., the full initial lump sum is allocated at the beginning of the period) that results in 7.41 total wealth over the full sample. This is consistent with the result from Figure 2. When k increases and τ decreases, we note that total wealth decreases eventually over the whole sample period (starting from 1994). This indicates that the investor waits too long before entering the stock market. We see that there is one BTD policy that results in relatively higher wealth of 7.57. This is attained for $k = 4$ and $\tau = -0.02, -0.01, 0.00$. Nonetheless, the main challenge is whether the investor can determine these values ex-ante. Second, if we consider the other parameters starting from 2000, we note that the investment policy $k = 4$ and $\tau = -0.02, -0.01, 0.00$ is not the most optimal. In fact, the investor can attain a higher terminal wealth by allocating a smaller fraction of her initial endowment over time and during more rare market selloffs. Third, we find a similar conclusion when we consider the other two periods. The main takeaway from Panel (a) from Table 1 is that BTD investment policies seem effective during periods that start with bull markets and are sensitive to selection of k and τ.

In terms of risk-adjusted performance, we report the Sortino ratio of the lump sum BTD investment policies in Panel (b) from Table 1. Specifically, we report differences in Sortino ratios between each policy and the passive one. Similar to the terminal wealth observations from Panel (a), Panel (b) shows that BTD investment polices result in higher risk-adjusted performance during periods that start with bear markets. For instance, since 2008 the investment policy of $k = 10$ and $\tau = -5\%$ not only increases terminal wealth by more than one million dollars, but it also enhances the Sortino ratio by 0.48. However, we observe that the risk-adjusted performance enhancement is less economically significant if one considers the testing period since 2000.

3.2 Monthly Inflows

We follow a similar discussion as before, however, with the focus on the monthly cash inflows. We consider the accumulation of wealth over time rather than a single endowment. At the end of each month, the investor allocates a single payment to a risky asset or defers investment for a "better" entry opportunity. The passive approach simply allocates to SPY as soon as these funds become available. The active investment policy follows the BTD approach, where the investor only allocates a unit of available funds to the risky asset if SPY drops below a certain level on a daily basis. Different from Section 3.1, we consider the problem from daily rather monthly monitoring.

In Table 2, we summarize the performance of 5 active BTD and 1 passive investment policies. The 5 active policies denote different levels of BTD depending on what level of τ determines the decision to BTD. Interestingly, Table 2 indicates that investors are better off accumulating terminal wealth using the passive investment policy regardless of the starting period. This implies that the opportunity cost of not investing the funds immediately could be relatively high. For example, investors who follow a BTD policy with $\tau = -0.02$ could have increased their terminal wealth by 14% if they have chosen the passive policy.

In terms of risk-adjusted performance, we observe that that the more conservative BTD policy ($\tau = -0.02$) leads to a higher Sortino ratio regardless of the time period under consideration. We also note that investors can attain a much higher Sortino ratio by allocating funds on a regular basis compared with the case of a single endowment. While it is a conventional wisdom that advocates allocating funds over time to your portfolio, investors with sufficiently high risk-aversion and/or short-term liquidity needs are likely to prefer a BTD strategy.

3.3 Maximum Drawdown

In the following discussion, we move to the MDD results. In Table 3, we summarize the results in line with Table 2. There are three main differences between this strategy and the previous one. First, the investor starts with 12 units of cash. Second, the MDD is computed using a rolling window of m-months daily returns. We set $m = 3, 6, 12$. For this reason, the first testing period starts from Jan 1995 to allow for 12 months of data to compute the first 12-months MDD. Third, when the current MDD is less than τ, the investor allocates all available cash holdings into the risky asset. We select three τ dip values of -20%, -10%, or -5% based on corresponding bear market, correction, and pullback heuristic technical thresholds commonly used by investors (Barron's (2019)).

A number of comments follow from Table 3. First, similar to the conclusion drawn above, a passive strategy leads to greater accumulated wealth over the testing period regardless of the specification. Second, we note that the -5% specification results in similar performance to the passive one. The reason for this is that such values for MDD are very common over the sample, leading to more frequent allocation to the risky asset as soon as funds become available. Third, we observe that the -20% MDD provides significant improvement over the passive strategy in terms of Sortino ratio and VaR. On the other hand, the -10% MDD does not provide much improvement in either metric. The evidence implies that the risk-adjusted return of the MDD policy is a U-shaped function of the threshold.[8]

4 Conclusion

Our analysis does not take into account liquidity shocks, but our risky asset SPY is one of the most traded securities in the world (Ben-David et al. (2016)) and is highly resistant to said shocks. Nevertheless, depending on investors' preference, BTD could provide utility in terms of liquidity provision over time. Overall, for mean-variance investors, BTD seems a natural way to maximize the risk-reward trade-off without the cumbersome task of estimation, evaluation, and prediction to form a strategic asset allocation. At the same time, if investors are concerned primarily with maximizing terminal wealth, a BTD policy may not be the optimal one to achieve this objective. This is subject to the initial market period state (i.e., bull or bear), which is generally beyond the control of an investor, as well as the determination of "dip" size τ. Determining the optimal τ ex-ante is similar to dealing with estimation risk in portfolio selection and, hence, does not guarantee ex-post optimality. Nonetheless, our results overall indicate that waiting for a bigger dip (e.g. a "bear market" decline of 20%) is better in terms of risk-adjusted performance. This is mostly evident for allocation of fixed inflows over time and beginning time periods that are associated with bear markets.

Dealing with estimation error requires advanced skills and perhaps proprietary data that is only accessible to sophisticated investors. Retail investors who try to emulate active management face major disadvantages in this area. Such disadvantage has important implications in terms of wealth distribution and, hence, inequality (Lei (2019)). More recently, fractional and commission-free brokers such as Robinhood have made it affordable for retail investors to participate in capital markets and inexpensively adjust portfolios. Whether it is a passive or BTD policy, our paper stresses the importance of continuous allocation to an investment account with or without the "dip". Finally, similar to the 1/N decision rule by DeMiguel et al. (2009), the question remains how inefficient BTD is as an investment policy compared to other backward/forward predictive models. We leave this for future research.

[8] In unreported results, we repeat the same analysis for a wide range of values of τ between -20% and 0% to confirm this U-shaped function observation.

References

Barron's, 2019. What's the difference between a correction and a bear market? https://www.barrons.com/articles/whats-the-difference-between-a-correction-and-a-bear-market-51553715750.

Basak, S., Chabakauri, G., 2010. Dynamic mean-variance asset allocation. The Review of Financial Studies 23, 2970–3016.

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., Vohs, K. D., 2001. Bad is stronger than good. Review of general psychology 5, 323–370.

Ben-David, I., Franzoni, F., Moussawi, R., 2016. Exchange traded funds (etfs). Tech. rep., National Bureau of Economic Research.

Best, M. J., Grauer, R. R., 1991. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. Review of Financial Studies 4, 315–342.

Clark, B., Feinstein, Z., Simaan, M., 2020. A machine learning efficient frontier. Operations Research Letters 48, 630–634.

Cookson, J. A., Niessner, M., 2020. Why don't we agree? evidence from a social network of investors. The Journal of Finance 75, 173–228.

Dancho, M., Vaughan, D., 2020. alphavantager: Lightweight R Interface to the Alpha Vantage API. R package version 0.1.2.

De Bondt, W. F., Thaler, R., 1985. Does the stock market overreact? The Journal of finance 40, 793–805.

DeMiguel, V., Garlappi, L., Uppal, R., 2009. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? Review of Financial Studies 22, 1915–1953.

DeMiguel, V., Martín-Utrera, A., Nogales, F. J., 2015. Parameter uncertainty in multiperiod portfolio optimization with transaction costs. Journal of Financial and Quantitative Analysis pp. 1443–1471.

Dimson, E., Marsh, P., Staunton, M., 2021. Credit Suisse Global Investment Returns Yearbook 2021 Summary Edition. https://www.credit-suisse.com/media/assets/corporate/docs/about-us/research/publications/credit-suisse-global-investment-returns-yearbook-2021-summary-edition.pdf.

Doran, J. S., Peterson, D. R., Wright, C., 2010. Confidence, opinions of market efficiency, and investment behavior of finance professors. Journal of Financial Markets 13, 174–195.

Fama, E. F., Fisher, L., Jensen, M. C., Roll, R., 1969. The adjustment of stock prices to new information. International economic review 10, 1–21.

Gentry, J., 2015. twitteR: R Based Twitter Client. R package version 1.1.9.

Goetzmann, W. N., Massa, M., 2002. Daily momentum and contrarian behavior of index fund investors. Journal of Financial and Quantitative Analysis 37, 375–389.

Golez, B., Koudijs, P., 2018. Four centuries of return predictability. Journal of Financial Economics 127, 248–263.

Jorion, P., 2007. Value at risk: the new benchmark for managing financial risk. The McGraw-Hill Companies, Inc.

Kim, W. C., Kim, J. H., Mulvey, J. M., Fabozzi, F. J., 2015. Focusing on the worst state for robust investing. International Review of Financial Analysis 39, 19–31.

Lei, X., 2019. Information and inequality. Journal of Economic Theory 184, 104937.

Loos, B., Previtero, A., Scheurle, S., Hackethal, A., 2020. Robo-advisers and investor behavior. Tech. rep., Working paper, University of Technology of Sidney.

Merton, R. C., 1973. An intertemporal capital asset pricing model. Econometrica: Journal of the Econometric Society pp. 867–887.

Michaud, R. O., 1989. The markowitz optimization enigma: is' optimized'optimal? Financial Analysts Journal 45, 31–42.

Nicolosi, G., Peng, L., Zhu, N., 2009. Do individual investors learn from their trading experience? Journal of Financial Markets 12, 317–336.

Sharpe, W. F., 2010. Adaptive asset allocation policies. Financial Analysts Journal 66, 45–59.

Sortino, F. A., Van Der Meer, R., 1991. Downside risk. Journal of portfolio Management 17, 27.

Wang, J., Taylor, N., 2018. A comparison of static and dynamic portfolio policies. International Review of Financial Analysis 55, 111–127.

Tables

Table 1: **Buy the Dip with Lump Sum**
Panel (a) below reports terminal wealth as a function of k (rows) and τ (columns) using lump sum BTD policies. The panel is split into four sub-panels. Each sub-panel denotes a different testing period. In all cases, the testing period dates between year y and 2020 (included). We choose $y \in \{1994, 2010\}$ to denote periods that began with bull markets and $y \in \{2000, 2008\}$ to indicate periods that began with bear markets. The first row $k = 1$ denotes the passive benchmark. Similar to Panel (a), Panel (b) reports the enhancement of risk-adjusted performance using a buy the dip (BTD) strategy over the passive one. Risk-adjusted performance is measured as Sortino ratio (Sortino and Van Der Meer (1991)).

Panel (a) Terminal Wealth

$k \setminus \tau$	-0.05	-0.04	-0.03	-0.02	-0.01	0.00	-0.05	-0.04	-0.03	-0.02	-0.01	0.00
			Since 1994						**Since 2010**			
1	7.41	7.41	7.41	7.41	7.41	7.41	3.47	3.47	3.47	3.47	3.47	3.47
2	5.65	7.57	7.40	7.40	7.40	7.40	3.49	3.49	3.54	3.54	3.54	3.54
3	4.99	7.59	7.51	7.51	7.51	7.51	3.57	3.57	3.53	3.53	3.53	3.53
4	4.38	7.05	7.54	7.57	7.57	7.57	3.47	3.59	3.57	3.57	3.57	3.57
5	4.02	6.57	7.12	7.56	7.56	7.56	3.47	3.51	3.59	3.59	3.59	3.59
6	3.81	6.12	6.71	7.56	7.56	7.56	3.38	3.50	3.52	3.49	3.47	3.53
7	3.69	5.77	6.31	7.26	7.26	7.33	3.17	3.42	3.51	3.45	3.40	3.45
8	3.62	5.37	5.98	6.95	6.95	7.10	2.95	3.23	3.44	3.45	3.35	3.38
9	3.58	5.06	5.63	6.69	6.69	6.84	2.79	3.09	3.32	3.40	3.33	3.33
10	3.58	4.83	5.33	6.41	6.41	6.62	2.65	2.92	3.21	3.33	3.34	3.30
			Since 2000						**Since 2008**			
1	2.43	2.43	2.43	2.43	2.43	2.43	2.69	2.69	2.69	2.69	2.69	2.69
2	2.50	2.50	2.50	2.50	2.50	2.50	2.79	2.79	2.79	2.79	2.79	2.79
3	2.52	2.52	2.50	2.50	2.54	2.54	2.90	2.90	2.90	2.84	2.84	2.84
4	2.58	2.58	2.51	2.51	2.53	2.53	3.04	3.04	3.04	2.92	2.88	2.88
5	2.66	2.66	2.56	2.56	2.53	2.53	3.26	3.26	3.26	3.03	2.93	2.93
6	2.74	2.74	2.63	2.63	2.53	2.53	3.44	3.44	3.44	3.21	2.98	2.98
7	2.81	2.81	2.70	2.70	2.53	2.53	3.61	3.61	3.61	3.37	3.05	3.05
8	2.90	2.90	2.77	2.74	2.56	2.54	3.81	3.81	3.81	3.53	3.18	3.18
9	2.96	2.96	2.86	2.80	2.61	2.56	3.78	3.78	3.79	3.72	3.31	3.31
10	3.04	3.04	2.91	2.87	2.66	2.58	3.77	3.77	3.76	3.72	3.44	3.44

Panel (b) Sortino Ratio

$k \setminus \tau$	-0.05	-0.04	-0.03	-0.02	-0.01	0.00	-0.05	-0.04	-0.03	-0.02	-0.01	0.00
			Since 1994						**Since 2010**			
1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	-0.10	0.01	-0.00	-0.00	-0.00	-0.00	0.04	0.04	0.01	0.01	0.01	0.01
3	-0.14	0.00	0.00	0.00	0.00	0.00	0.06	0.06	0.04	0.04	0.04	0.04
4	-0.19	-0.03	0.00	0.00	0.00	0.00	0.05	0.07	0.06	0.06	0.06	0.06
5	-0.22	-0.05	-0.02	0.00	0.00	0.00	0.06	0.05	0.07	0.07	0.07	0.07
6	-0.24	-0.08	-0.04	-0.00	-0.00	-0.00	0.04	0.06	0.06	0.04	0.03	0.05
7	-0.25	-0.10	-0.07	-0.02	-0.02	-0.01	-0.03	0.04	0.07	0.03	-0.00	0.02
8	-0.25	-0.12	-0.08	-0.03	-0.03	-0.02	-0.10	-0.01	0.06	0.04	-0.03	-0.01
9	-0.26	-0.14	-0.10	-0.05	-0.05	-0.04	-0.15	-0.06	0.02	0.03	-0.03	-0.03
10	-0.25	-0.16	-0.12	-0.06	-0.06	-0.05	-0.19	-0.12	-0.02	0.01	-0.02	-0.04
			Since 2000						**Since 2008**			
1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.02	0.02	0.02
3	0.01	0.01	0.01	0.01	0.02	0.02	0.05	0.05	0.05	0.03	0.03	0.03
4	0.02	0.02	0.01	0.01	0.01	0.01	0.10	0.10	0.10	0.05	0.04	0.04
5	0.04	0.04	0.02	0.02	0.01	0.01	0.21	0.21	0.21	0.09	0.05	0.05
6	0.05	0.05	0.03	0.03	0.01	0.01	0.28	0.28	0.28	0.18	0.06	0.06
7	0.06	0.06	0.04	0.04	0.01	0.01	0.34	0.34	0.34	0.25	0.09	0.09
8	0.08	0.08	0.06	0.05	0.01	0.01	0.43	0.43	0.43	0.31	0.15	0.15
9	0.09	0.09	0.07	0.06	0.03	0.01	0.46	0.46	0.45	0.39	0.21	0.21
10	0.10	0.10	0.08	0.07	0.04	0.02	0.48	0.48	0.47	0.42	0.27	0.27

Table 2: **Buy the Dip with Monthly Inflows**

The table below reports differences in performance of a buy the dip (BTD) strategy with monthly inflows using different τ dip values. The first panel corresponds to terminal wealth, whereas the other two panels report risk-adjusted performance using the Sortino ratio (Sortino and Van Der Meer (1991)) and Value-at-Risk (VaR). VaR (reported in percentages) is computed with respect to Jorion (2007) as the average monthly return minus the 1% bottom percentile. The results are reported as a function of τ (columns). In each panel, the rows denote a different testing period. In all cases, the testing period dates between year y and 2020 (included). We choose $y \in \{1994, 2010\}$ to denote periods that began with bull markets and $y \in \{2000, 2008\}$ to denote periods that began with bear markets. The column "Bull" indicates 1 (0) whether the start of the backtesting period corresponds to a bull (bear) market period. The last column corresponds to the passive strategy in which the investor allocates funds to the risky asset SPY as soon as they become available at the beginning of the month.

Bull	Year Start	$\tau = -0.020$	$\tau = -0.015$	$\tau = -0.010$	$\tau = -0.005$	$\tau = 0.000$	Passive
			Terminal Wealth				
1	1994	1246.44	1309.07	1411.40	1430.41	1432.41	1434.59
1	2010	246.22	285.13	301.79	304.01	304.66	305.27
0	2000	824.45	817.78	835.72	837.33	838.02	839.39
0	2008	356.54	394.67	410.95	412.73	413.08	413.84
			Sortino Ratio				
1	1994	2.83	2.56	2.56	2.57	2.54	2.54
1	2010	7.81	6.19	5.79	5.82	5.82	5.83
0	2000	3.73	3.31	3.19	3.22	3.22	3.23
0	2008	6.63	5.78	5.47	5.55	5.54	5.53
			Value-at-Risk				
1	1994	11.58	12.68	12.77	12.78	12.77	12.77
1	2010	9.80	11.80	12.86	12.87	12.88	12.89
0	2000	11.14	11.90	12.21	12.23	12.23	12.24
0	2008	10.31	11.60	12.23	12.26	12.26	12.26

Table 3: **Buy the Dip with Maximum Drawdown**
The table below reports the performance summary of a buy the dip (BTD) strategy based on maximum drawdown (MDD). The first panel corresponds to terminal wealth, whereas the other two panels report the risk-adjusted performance using Sortino ratio (Sortino and Van Der Meer (1991)) and Value-at-Risk (VaR). VaR (reported in percentages) is computed with respect to Jorion (2007) as average monthly return minus the 1% bottom percentile. The results are reported as a function of τ (columns), which is set to be either -20%, -10%, or -5%. We consider three different MDD measures. Each is computed on m months of rolling daily window, where m is set to either 3, 6, or 12. If the m-months MDD is less than τ on a given day, the investor allocates all of her cash into the risky asset. In each panel, the rows denote a different testing period. In all cases, the testing period dates between year y and 2020 (included). We choose $y \in \{1995, 2010\}$ to denote periods that began with bull markets and $y \in \{2000, 2008\}$ to denote periods that began with bear markets. The column Bull indicates 1 whether the start of the backtesting period corresponds to a bull period. The Passive columns denote the allocation policy that allocates funds to the risky asset as soon as the funds become available.

		3-months MDD			**6-months MDD**			**12-months MDD**			**Passive**
Bull	Year	-20%	-10%	-5%	-20%	-10%	-5%	-20%	-10%	-5%	
		Panel (a) Terminal Wealth									
1	1995	1178.11	1254.83	1403.82	1075.69	1267.06	1413.92	1072.06	1280.09	1419.93	1419.99
1	2010	215.71	332.42	349.60	216.17	338.87	349.98	250.11	343.13	350.01	350.42
0	2000	846.69	826.79	880.25	797.10	835.58	881.46	793.46	851.08	880.89	881.30
0	2008	349.39	430.27	450.26	349.78	438.17	450.38	359.35	445.29	450.42	450.83
		Panel (b) Sortino Ratio									
1	1995	1.75	1.55	1.59	1.60	1.55	1.59	1.55	1.55	1.60	1.60
1	2010	7.44	2.90	3.21	7.35	2.95	3.22	7.91	3.07	3.22	3.22
0	2000	2.24	1.81	1.89	1.97	1.84	1.90	1.89	1.86	1.90	1.90
0	2008	4.74	2.78	2.92	3.23	2.82	2.92	3.32	2.84	2.92	2.92
		Panel (c) Value-at-Risk									
1	1995	10.79	11.86	11.97	10.59	11.87	11.98	10.82	11.87	11.98	11.98
1	2010	3.44	10.34	10.68	3.46	10.44	10.69	4.29	10.64	10.69	10.69
0	2000	9.70	11.01	11.12	9.84	11.03	11.12	10.26	11.06	11.12	11.12
0	2008	6.10	11.77	11.86	7.30	11.81	11.86	7.68	11.87	11.86	11.86