Buy the Dip?[*]

Stefano Bonini[†] [‡] Thomas Shohfi[§¶] Majeed Simaan[†]

sbonini@stevens.edu shohfit@sec.gov msimaan@stevens.edu

January 13, 2023

[*]We are grateful for valuable comments from Steve LeCompte, Neev Vora, Steven Pav, George Calhoun, and Marko Gränitz. The statistical computation of this paper is conducted using R (Team R Core and others). The source code is available at https://github.com/simaan84/btd. Building on our empirical analysis, interested readers/users can evaluate different BTD strategies over different periods and assets using the dashboard created by the authors of this paper. The dashboard is available at http://btd.today/. Our analysis relies on multiple open-source libraries thanks to the many contributing members of the R community such as Jeffrey A. Ryan and Joshua M. Ulrich as well as Brian G. Peterson and Peter Carl. All errors remain our responsibility. The authors are providing information but make no direct fiduciary recommendations on its implementation. This paper is to be used by an investor by their own choice and merits. One should consult a financial advisor if needed.

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

[‡]Corresponding author.

[§]Division of Economic and Risk Analysis (DERA), U.S. Securities and Exchange Commission, 100 F Street NE, Washington, DC 20549, USA.

Buy the Dip?*

January 13, 2023

Abstract

We study the fundamental properties of the ‘Buy the dip” (BTD) investment heuristic. Looking into cash holdings versus a stock market exchange-traded fund, we find that BTD does not necessarily maximize investors’ real terminal wealth and is sensitive to market conditions at the beginning year of investment. While under certain conditions, BTD may improve risk-adjusted performance over a passive investment policy or a classical Dollar Cost Averaging approach, its optimality is subject to estimation risk. Given the vast popularity of BTD, our results have important implications for asset managers and retail investors alike.

Keywords: *Asset Allocation, Model Risk, Market Timing, Dollar Cost Averaging*

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

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

> No matter how pessimistic you are (and everybody else is), and no matter how bad the financial and world news is, you must not interrupt the plan or you will lose the important benefit of insuring that you buy at least some of your shares after a sharp market decline.
>
> *Malkiel (1999)*

1 Introduction

Buy the dip (henceforth BTD) is a heuristic well known among both professional and retail investors. According to regulatory filings, billionaire investor Warren Buffett has utilized the market selloff during the first quarter of 2022 to "**buy-the-dip** and add several new major positions", corresponding to tens of billions of dollars investment (Klebnikov, 2022). The popularity of the BTD heuristic is further supported by the sharp increase in public interest measured by search popularity, as shown in Figure 1, and its frequent reference in financial news and asset managers communications, as captured by over 4.5 million news articles included in Factiva that mention BTD investment strategies in the 2008-2022 period.

INSERT FIGURE 1 HERE

A BTD strategy can be summarized as the purchase of an asset after a sharp decline in price and in anticipation of a reversal. Conceptually, BTD can be interpreted as an extension of the popular dollar cost averaging approach that is ubiquitously implemented by asset managers. However, while the investment features of DCA have been extensively studied (Kirkby et al. (2020), Brennan et al. (2005), Constantinides (1979),) little is known about BTD strategies. In particular a number of questions are unanswered: What constitutes a dip? What is the unconditional performance of such a strategy? What is its performance conditional on market cycles? In this paper we attempt at filling this gap providing also preliminary (indirect) evidence on its adoption by market participants.

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 (see, e.g., Fisher and Statman (2006)). 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, sell-offs could also reflect the persistence of bad news in the market. As a result, accumulating cash during growth periods to be invested following significant market drops may not reflect the most optimal investment policy. However, behavioral studies recognize that bad news has a stronger effect than good news (Baumeister et al. (2001)). Hence, if market participants overreact to bad news, BTD could potentially lead to a correction in the market and constitute an optimal investment.

According to Sharpe (2010), a strategic asset allocation denotes a process with well-defined steps that define a specific allocation of capital to the underlying assets, i.e., the policy portfolio. Following this approach, in this paper, we identify 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 relative performance of a BTD-based investment in the market portfolio (proxied by the SPY ETF), 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. We consider two main endowment setups 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.

We obtain several important results. First, we highlight that a BTD policy may enhance the terminal wealth of the investor over time subject 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 the start date of the strategy (Dimson et al., 2021). However, there is much potential improvement of risk-adjusted performance when investors use maximum drawdown (MDD) to measure potential dip.[1]

Second, investors face uncertainty when it comes to anticipating all states of the world that are relevant for the investment policy problem (Golosnoy and Okhrin, 2008). Such uncertainty raises concerns regarding whether a pre-determined policy is optimal. For instance, mean-variance optimization is designed to improve investors' risk-adjusted returns. However, 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 that improve risk-adjusted returns over time.[2] This result is particularly apparent in the case of monthly cash inflows. While investors may not maximize their terminal wealth using BTD compared to a passive policy, they can achieve a higher risk-adjusted reward with lower downside risk by "buying the dip".

Third, a BTD strategy is inherently contrarian, which is more common than a momentum strategy among retail investors. These long-term contrarian retail investors have asymmetric market-timing preferences as they are more likely to "buy the dips" instead of "selling the rips", a similar but opposite strategy entailing selling at a peak (Goetzmann and Massa,

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

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

2002).[3] To account for this characteristic, we consider a simulation study to control for different market conditions. The experiment is motivated by the return models studied by Kaminski and Lo (2014). Specifically, we consider an autoregressive model of order 1, i.e., AR(1). This provides a parsimonious way to control for different market regimes. We calibrate the AR(1) model using the first two moments of daily returns of the SPY ETF. Based on which, we simulate five years of daily returns using 11 regimes, capturing random walk, mean-reversion, and momentum. Given the simulated data, we rerun both BTD and passive investment policies and report several performance metrics. Not surprisingly, we find that the BTD policy dominates a passive strategy in terms of risk-adjusted returns and terminal wealth in the experiments associated with large negative serial correlation, i.e., market reversal. More interestingly, we find that BTD consistently results in lower position volatility and Value-at-Risk (VaR). However, this comes with a large opportunity cost. In cases of high market momentum, BTD exhibits large downside in terms wealth accumulation, even though it succeeds in lowering portfolio volatility and VaR relative to a passive strategy. Indeed, according to the panel regression analysis utilizing the simulated data, we find that the relationship between risk and wealth using BTD is reduced by 30% after taking into account market regime fixed effects.

Fourth, while our paper considers 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 on 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. However, this result is also subject to estimation risk as investors need to determine the "optimal" parameters ex-ante. After all, investors with high risk tolerance are better off pursuing a passive investment strategy as their utility is less sensitive to portfolio' volatility smoothing

[3]Contrary to Goetzmann and Massa (2002), Hoffmann and Shefrin (2014) find that retail investors are less inclined to bet on reversals. Their empirical study, however, utilizes a smaller sample based on Dutch discount brokerage clients covering survey responses of 5,500 retail investors.

resulting from BTD.

Finally, our paper has important implications for robo-advisers and their impact on retail investors (Loos et al., 2020). While BTD is subject to estimation error, it provides better transparency into decision-making than predictive models that rely on historical data (Clark et al., 2020). Given the challenges in the academic literature in finding the optimal portfolio (DeMiguel et al., 2009), a BTD policy provides a simple heuristic for investors to allocate their wealth over time with lower complexity. The remainder of this manuscript proceeds as follows. In Section 2, we discuss the related literature. Section 3 describes the main investigation behind our analysis, covering different BTD trading strategies. In Section 4, we discuss the main results and findings of the paper based on empirical data. We devote Section 5 to Monte Carlo simulation studying the appeal of BTD under different market conditions/regimes. Section 6 concludes.

2 Related Literature

As illustrated by the opening quote by Malkiel (1999), investors have long followed different heuristic investment approaches, one of which is known as dollar-cost averaging (DCA). Such an investment strategy can be traced back to the 1940s (Meir, 1995). By definition, a DCA strategy averages a lump sum over the time horizon using a predetermined plan rather than investing the total amount at the beginning of the period. Its performance however is debated. Constantinides (1979) shows that the DCA strategy is sub-optimal from a theoretical point of view. Differently, Brennan et al. (2005) find that "individual investors who follow [DCA] strategy in purchasing individual stocks to add to an existing portfolio are better off than if they followed the 'rational' strategies traditionally recommended by academics."

While BTD shows some similarities to DCA, it differs from DCA in multiple ways. In a DCA plan, the allocation/averaging is predetermined ex-ante. Differently, BTD is a contingent policy depending on the measurement/identification of the dip, e.g., the threshold that investors view as a dip as well as the frequency used to measure the dip size. In either case, investors need to determine the plan ex-ante. In the case of BTD, this plan corresponds to the measurement and identification of the dip size. Furthermore, while DCA is well suited for lump sum investment decisions, we extend the analysis of BTD properties to a more general case involving the allocation of monthly cash flows as they become available.

In testing BTD strategies (and *a fortiori* also DCA) we face a severe methodological problem. Rational strategies, which are traditionally recommended by academics (Brennan et al., 2005), ignore estimation risk in general. Estimation risk corresponds to the category of model risk and constitutes a major challenge in constructing optimal risk-reward strategies (Best and Grauer, 1991). A potential remedy in the academic literature is utilizing shrinkage techniques (Ledoit and Wolf, 2003; DeMiguel et al., 2009; Kan and Zhou, 2007). The main idea behind shrinkage is finding an optimal bias-variance trade-off. On the one hand, by incorporating a heuristic into the optimized portfolio, investors introduce a bias into the "supposedly optimal" portfolio, which could cause a systematic shift from maximum utility. On the other hand, such a prior helps mitigate portfolio variability and, hence, introduces greater stability into asset allocation, translating to lower transaction costs. For instance, a heuristic approach that allocates equally among risky assets has been known to outperform common decision rules studied in the literature (DeMiguel et al., 2009). In this regards, a BTD policy can be considered as a bias introduced into the decision-making process. Not surprisingly, such a decision does reduce portfolio variance and, hence, enhances risk-adjusted returns, implying higher utility for risk-averse investors. Nonetheless, to implement the BTD policy, investors need to determine the dip size, which corresponds to estimation risk. From a theoretical point of view, estimation risk raises concerns in determining optimal out-of-sample portfolios (Simaan and Simaan, 2019). While the DCA plan requires certain specifications to implement, it is subject to less estimation risk relative to BTD. Specifically, BTD is more challenging and is subject to estimation risk, raising questions about its appeal from an ex-ante point of view. While our results resonate with those from Brennan et al. (2005), the challenge with BTD is determining the "correct" specifications to implement an out-of-sample optimal policy.

3 Data and Methodology

3.1 Data

While a BTD strategy could be implemented with any asset, for generality we focus on its application to an investment in the market index, captured by the SPY, the largest and oldest exchange-traded fund (ETF) tracking the S&P 500 index. The data set includes two time series. The first is the daily return of the SPY from Jan 1994 to Dec 2020. The second 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

$$\tilde{R}_{t+1} = \frac{1 + R_{t+1}}{1 + i_{t+1}} - 1 \tag{3.1}$$

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 consideration the inflation-adjusted return with respect to Equation (3.1). Adjusted for inflation, Panel (b) indicates that a passive investment in SPY would have grown a \$1 million investment to \$7.4 million.

INSERT FIGURE 2 HERE

3.2 Portfolio Selection

3.2.1 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 notation of the inflation rate and inflation-adjusted return, investor wealth over the next period is defined as

$$w_{t+1} = \theta_t w_t (1 + \tilde{R}_{t+1}) + (1 - \theta_{t+1}) w_t \frac{1}{1 + i_{t+1}} \tag{3.2}$$

The above can be rewritten as

$$w_{t+1} = \frac{w_t}{1 + i_{t+1}} \left[1 + \theta_t R_{t+1}\right] \tag{3.3}$$

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

$$\theta_t^* = \frac{1}{A} \frac{\mu_{t+1|t}}{\sigma^2_{t+1|t}} \tag{3.4}$$

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 (3.4) 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.[4] It is well-documented that estimation error in mean-variance portfolio selection is severe and leads to poor out-of-sample performance.[5] In this respect a BTD heuristic may offer a simpler decision rule for investors at large without requiring complex portfolio optimization, similar in spirit to the outcomes of naive passive strategies highlighted by DeMiguel et al. (2009).

3.2.2 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 investors 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 a given time period. The measure provides

[4] 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.

[5] See, e.g., Michaud (1989); Best and Grauer (1991); Sheedy et al. (1999); DeMiguel et al. (2009) and references therein.

a risk-reward assessment of the underlying investment policy. Risk is measured by 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 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.

3.3 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.

3.3.1 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%.

3.3.2 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.[6] 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 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 only if its daily return drops below τ. This is subject to the constraint that her outstanding funds are greater than zero. Our implementation sets $\tau = \{-0.020, -0.015, -0.010, -0.005, 0.000\}$.

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

3.3.3 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.

3.3.4 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 (i.e., immediately investing each cash flow upon receipt in SPY regardless of market conditions).

4 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.

4.1 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.

INSERT FIGURE 3 HERE

In Panel (a) from Table 1, we report terminal wealth for different policies discussed in Section 3.3. 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 bear markets and are sensitive to selection of k and τ.

INSERT TABLE 1 HERE

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.

INSERT FIGURE 4 HERE

Figure 4 provides a visual illustration of the results from Panel (a) of Table 1. Overall, the results are consistent with the insights derived from the table. The first column on the left-hand side of each panel denotes the passive strategy, whereas the bottom right corner represents the more conservative BTD investment policy (i.e., large k and small τ). The red (blue) colors of the figure denote high (low) levels of terminal wealth. Depending on the testing period, we observe that moving diagonally from the top left to the bottom right could potentially enhance an investor's terminal wealth.

4.2 Monthly Inflows

We follow a similar discussion as before, however, with the focus on monthly cash inflows. We consider accumulation of wealth over time rather than via 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 4.1, 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 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 had chosen the passive policy.

INSERT TABLE 2 HERE

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 conventional wisdom advocates allocating funds over time to one's portfolio, investors with sufficiently high risk-aversion and/or short-term liquidity needs are likely to prefer a BTD strategy.

4.3 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)).

INSERT TABLE 3 HERE

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 a 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 a 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.[7]

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

4.4 Information Flow

In market microstructure theory, price and volume changes are related to the arrival of information to the market (see, e.g., Li and Wu (2006)). In the following discussion, we approximate information flow as a relative change in volume multiplied by the return on the SPY. This measure gauges whether the number of shares being traded is positively associated with the market movement. The purpose of this empirical exercise is to investigate the determinants of the BTD strategy and better understand whether market participants purchase a higher number of shares surrounding dip periods.

We follow a similar approach used in the previous section, where we determine the dip using MDD based on a three-month rolling window of daily returns. For each quarter in the data, we measure the maximum MDD for that quarter. This results in 107 MDDs, from which we choose the top 50% to denote the more significant events, corresponding to 53 quarters. From these quarters, we identify the exact dates associated with such large drops, which we refer to as events. In total, we have 306 daily events where the MDD was above the median quarterly peak. For each event, we compute the cumulative 21-day returns before and after the event. Finally, we use a t-test to determine whether the difference between the two is statistically significant.

INSERT TABLE 4 HERE

Table 4 summarizes the results of the above test. The first three columns correspond to the full sample period. This specification is added as a robustness check to validate whether such difference happens across the whole sample or is only evident during the event days. Overall, we find no statistical evidence to support the argument that there is a change in volume and return on a day-to-day basis. However, considering the sub-sample associated with the dip events, we find strong statistical evidence of increased information flow. This result suggests that market participants overall recognize just events, where they view the market on a discount and allocate more capital after major drops. These findings provide an intuitive explanation of why the BTD strategy may actually work. After all, if all market participants are anticipating such events, then BTD becomes a self-fulfilling prophecy and, hence, results in higher risk-adjusted returns.

5 Additional Results

In this section, we consider numerical analysis to gain a deeper understanding behind the appeal of BTD from risk-adjusted returns. Specifically, we simulate market returns under different conditions and follow a similar investigation from Section 3.

5.1 Market Returns

We consider a number of return-generating processes for market returns, including random walk, mean-reversion, and momentum. These specifications are consistent with the data-generating functions studied by Kaminski and Lo (2014).

Definition 1 *Let R_t denote the return on the risky asset over a sequence of periods $t = 1, ..., T$. We define the return-generating process of the market return R_t as follows:*

1. ***Random Walk**. Under the random walk hypothesis, the return process $\{R_t\}_{t=1}^{T}$ follows a random walk such that*

$$R_t \overset{IID}{\sim} N(\mu, \sigma^2) \tag{5.1}$$

2. ***Mean Reversion/Momentum**. A non-random walk can be modeled using an autoregressive model of order 1, i.e., AR(1), where*

$$R_t = \alpha + \rho R_{t-1} + \epsilon_t \tag{5.2}$$

with $\epsilon_t \overset{IID}{\sim} N(0, \sigma_\epsilon^2)$ for $t = 1, .., T$ denoting white noise.

Note that the second specification requires that $\rho \in (-1, 1)$ to assure stationarity. Additionally, note that the random walk can be retrieved from this process by setting $\rho = 0$. On the other hand, when $\rho > 0$, we expect previous returns to persist over the next period, capturing the idea of momentum. In a similar manner, when $\rho < 0$, we expect a contrarian process where the next period return is negatively related to the previous return, implying a reversal in the market.

5.2 Calibration

In order to implement the simulation analysis based on the above data-generating functions, we need to calibrate the parameters. For instance, in the case of the random walk, we only need to estimate μ and σ_ϵ to reflect daily mean returns and volatility. For the AR(1) model, on the other hand, we can assess the sensitivity of the model with respect to different specifications of ρ. However, to follow a consistent analysis, we calibrate the parameters such that both the expected return and the volatility of the process are equal to the calibrated μ and σ_ϵ from the random walk.

For a stationarity time series, recall that

$$\mu = \alpha + \rho\mu \rightarrow \mu = \frac{\alpha}{1 - \rho} \tag{5.3}$$

Hence, given the calibrated parameters from the random walk and an arbitrary value of ρ, we can calibrate the parameter α using the above condition. The same reasoning follows for volatility, where we have

$$\sigma^2 = \frac{\sigma_\epsilon^2}{1 - \rho^2}. \tag{5.4}$$

In this case, for a calibrated σ and an arbitrary value of ρ, we can determine the value of σ_ϵ that satisfies the above condition. Overall, we observe that we can calibrate the parameters of the AR(1) return process using arbitrary values of ρ.

5.3 Simulation

We consider arbitrary values of $\rho \in \{-0.5, -0.4, \cdots, 0, \cdots, 0.4, 0.5\}$ to simulate the above return process. Note that we can retrieve the random walk when $\rho = 0$ and reversal (momentum) when $\rho < 0$ ($\rho > 0$). For a given ρ, we perform the following steps:

1. Choose α and σ_ϵ such that both (5.3) and (5.4) are satisfied.

2. Generate five years of daily returns using the data process.

3. Implement the same BTD experiment from Section 3.3.2 using the simulated data and evaluate performance. In all cases, set $\tau = -0.01$, i.e., allocate one unit of cash to the risky asset each time the daily return drops below -1%.

4. Repeat the last two steps 1000 times.

We run the above experiment for $\rho \in \{-0.5, -0.4, \cdots, 0, \cdots, 0.4, 0.5\}$. From each experiment, we compute similar performance metrics as before such as Sortino ratio, VaR, terminal wealth, and position volatility. After dropping extreme/missing values, the experiment results in 10,963 observations, which closely correspond to 11×1000 observations. The data denotes a ρ-trial panel, which provides interesting insights about the performance of BTD relative to a passive strategy based on different return model specifications. Specifically, from each experiment, we compute the difference in metrics between BTD and passive strategies. Based on these observations, we summarize our main simulation results in the following discussion.

5.4 Baseline Results

In Table 5, we report the median difference in performance with respect to each level of ρ. Specifically, for a performance metric X, the difference in performance is computed as

$$\Delta X = X_{BTD} - X_{Passive}$$

with X_{BTD} and $X_{Passive}$ denoting the performance metric of BTD and passive strategies, respectively. For each value of ρ, we have roughly 1000 observations. Based on which, we compute the median of ΔX for different performance metrics. We report the results in a descending order with respect to ρ.

INSERT TABLE 5 HERE

A couple of comments follow from Table 5. First, in terms of risk-adjusted performance, we observe a marginal improvement of BTD over a passive strategy for small values of ρ. This observation is consistent with the enhancement in wealth using BTD. However, for higher values of ρ, especially when the market is in momentum, we observe that BTD underperforms consistently relative to the passive strategy. This result should not be surprising. After all, BTD is a contrarian strategy.

Second, and perhaps more interesting, we observe that BTD results in lower risk regardless of market conditions. This is evident from both the volatility and VaR measures. Additionally, we note that the advantage of BTD in terms of risk reduction strategy increases as we increase the value of ρ. A potential explanation for this result is the following. Since the data is generated using a positive mean, on average the market will go up. In a market where momentum is likely to persist, there are fewer opportunities to buy the dip, such that the investor ends up holding more cash and staying out of the market. While this results in lower returns and wealth, this also implies lower volatility and downside risk.

To gain a deeper perspective on the simulation results, we plot the distribution of the performance metrics for the 11 values of ρ in Figure 5. Panel (a), for instance, demonstrates the difference in position volatility between BTD and passive strategies. The bottom distribution of the panel denotes the case of reversal with $\rho = -0.5$, whereas the top one denotes momentum with $\rho = 0.5$. The distributions are ranked from top to bottom in a descending order with respect to ρ. The middle distribution denotes the case of a random walk, i.e., $\rho = 0$. For each distribution, we trim the plot at 1% and 99% to focus on the majority of the data and less on extreme outliers. In all cases, we add a vertical dashed line to denote the zero level.

INSERT FIGURE 5 HERE

A number of interesting insights follow from Figure 5. First, the plot confirms the insights from Table 5, where BTD's advantage over the passive strategy in terms of risk reduction is evident in Panels (a) and (c). We observe this from the fact that most of the distribution is located to the left hand side of the dashed horizontal line regardless of market conditions. Nonetheless, we observe that there is greater variability in this advantage as ρ increases. Since the plots are trimmed at 1% and 99%, we observe that BTD has more "good" extreme cases of reducing risk, i.e., more left-tail chance than right-tail. A similar observation follows with respect to VaR.

INSERT TABLE 5 HERE

Second, in terms of of risk-adjusted returns (Sortino ratio) and wealth (Panels (b) and (d)), the advantage of BTD is only evident in cases of high reversal, i.e., negative values of ρ. Additionally, we note that the difference in wealth exhibits much greater variability relative to the Sortino ratio. A possible explanation for this result is the risk reduction evidence of BTD. Regardless of the value of ρ, we note that BTD does mitigate the risk of the portfolio across all distributions, which results in a lower dispersion in the Sortino ratio. Nonetheless, we also note that there are more extreme cases when evaluating both metrics.

For instance, the left-side in Panel (d) is too wide, indicating a higher probability of extreme underperfomance in terms of wealth accumulation. On the other hand, the result for Sortino ratio seems roughly the same from both sides.

5.5 Effect of AR(1)

Earlier studies reveal that the random walk process is rejected in practice and present evidence of mean reversion in stock returns (De Bondt and Thaler, 1985). To better understand the relationship between the simulation results and our empirical analysis, we look into serial correlation of the SPY daily return using the full sample period. We consider a rolling window estimation that is rolled on a daily basis and utilizes a sample window of 252 or 504 trading days, reflecting one year and two years, respectively. This coefficient corresponds to the estimated ρ parameter over time. We illustrate the empirical evidence in Figure 6.

INSERT FIGURE 6 HERE

On average, the estimated ρ is negative, roughly -6%. While we observe large variability, more than 75% of the time the estimate is negative. A major drop in the estimate occurs during early 2020, which is associated with the COVID-19 selloff and followed by one of the quickest recoveries. In terms of BTD implications, it is therefore important to relate our findings to market conditions in terms of reversal versus momentum.

In what follows, we utilize the simulation results to estimate the effect of the ρ parameter on the difference between BTD and passive strategies in terms of risk-reward trade-off. In order to do so, we run a number of regressions using the simulation panel. We set the dependent variable to denote either the differential in mean return or differential in terminal wealth. As the independent variable, we consider the differential in position volatility. In other words, we are interested in addressing the question of how much investors can enhance their risk-reward trade-off under different market conditions. To control for different market regimes, we also utilize a panel regression with ρ as fixed effects.

INSERT TABLE 6 HERE

We report the regression results in Table 6. Overall, we find significant evidence in terms of a positive risk-reward trade-off. A larger increase risk-taking of BTD commensurates with both higher returns and higher terminal wealth. For mean return, we observe that the effect becomes relatively stronger when we control for fixed effects. On the other hand, we find that the effect becomes weaker for terminal wealth after controlling for fixed effects. Overall, this result sheds important light on how the additional risk-taking via BTD compensates in terms of reward.

INSERT FIGURE 7 HERE

In Figure 7, we illustrate the coefficients of the fixed effects from the panel regressions reported in Table 6. Consistent with the previous insights, we observe that the fixed effects are associated with a positive (negative) trend when the dependent variable is mean return (terminal wealth). In regimes with high ρ, investors end up with a greater amount of cash, reducing risk. When market losses persist, the BTD strategy enhances the risk-reward trade-off due to lower risk-taking. On the other hand, when market gains persist, the lower participation in the stock market via BTD enhances overall the risk-reward trade-off. Not surprisingly, however, we observe that the fixed effects for the terminal wealth decrease with ρ, which is consistent with the previous reasoning.

6 Conclusion

Buy the dip (BTD) has emerged as one of the most popular heuristics among both professional and retail investors due to its intuitive appeal and its similarity with Dollar Cost Averaging strategies. Surprisingly though, its properties and the conditions (if any) under which it may generate positive absolute and risk-adjusted returns are largely unknown. In this paper we provide a comprehensive analysis of the performance of BTD conditional on: the size of the dip, the market conditions at the start of the strategy and the selected benchmark. We focus on the SPY as the investable risky asset due its exceptional liquidity (Ben-David et al. (2016)) in order to minimize confounding effects coming from exogenous liquyidity shocks.

We obtain a number of results. First, 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 (Nam and Branch (1994)). Second, BTD returns (similarly to DCA) are subject to the initial market period state (i.e., bull or bear) - which is generally exogenously set for investors - as well as the selection of the "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 indicate that investment rules targeting larger dips (e.g. a "bear market" decline of 20%) yield better risk-adjusted performance than higher frequencies approaches based on smaller dip sizes. This is mostly evident for allocation of fixed inflows over time and beginning time periods that are associated with bear markets. Third, a BTD strategy is weakly under-performing an optimally balanced benchmark. However dealing with estimation error requires advanced skills and access to proprietary data that is only accessible to sophisticated investors (Farboodi et al. (2021)). As highlighted by Lei (2019), retail investors who try to emulate active management are therefore penalized with important implications in terms of wealth distribution and, hence, inequality. Differently, a simple BTD policy, especially with the advent of fractional and commission-free brokers such as Robinhood, allows retal investors to target similar returns at low cost and without requiring access to high-level skills and computational power. Crucially though, results also stress 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.

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. and G. Chabakauri (2010). Dynamic mean-variance asset allocation. *The Review of Financial Studies 23*(8), 2970–3016.

Baumeister, R. F., E. Bratslavsky, C. Finkenauer, and K. D. Vohs (2001). Bad is stronger than good. *Review of general psychology 5*(4), 323–370.

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

Best, M. J. and R. R. Grauer (1991). On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. *Review of Financial Studies 4*(2), 315–342.

Brennan, M. J., F. Li, and W. N. Torous (2005). Dollar cost averaging. *Review of Finance 9*(4), 509–535.

Clark, B., Z. Feinstein, and M. Simaan (2020). A machine learning efficient frontier. *Operations Research Letters 48*(5), 630–634.

Constantinides, G. M. (1979). A note on the suboptimality of dollar-cost averaging as an investment policy. *Journal of Financial and Quantitative Analysis 14*(2), 443–450.

De Bondt, W. F. and R. Thaler (1985). Does the stock market overreact? *The Journal of finance 40*(3), 793–805.

DeMiguel, V., L. Garlappi, F. J. Nogales, and R. Uppal (2009). A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management science 55*(5), 798–812.

DeMiguel, V., L. Garlappi, and R. Uppal (2009). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *Review of Financial Studies 22*(5), 1915–1953.

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

Dimson, E., P. Marsh, and M. Staunton (2021, March). 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.

Farboodi, M., D. Singal, L. Veldkamp, and V. Venkateswaran (2021). Valuing financial data. *Available at SSRN 3947931*.

Fisher, K. L. and M. Statman (2006). Market timing in regressions and reality. *Journal of Financial Research 29*(3), 293–304.

Goetzmann, W. N. and M. Massa (2002). Daily momentum and contrarian behavior of index fund investors. *Journal of Financial and Quantitative Analysis 37*(3), 375–389.

Golosnoy, V. and Y. Okhrin (2008). General uncertainty in portfolio selection: A case-based decision approach. *Journal of Economic Behavior & Organization 67*(3-4), 718–734.

Hoffmann, A. O. and H. Shefrin (2014). Technical analysis and individual investors. *Journal of Economic Behavior & Organization 107*, 487–511.

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

Kaminski, K. M. and A. W. Lo (2014). When do stop-loss rules stop losses? *Journal of Financial Markets 18*, 234–254.

Kan, R. and G. Zhou (2007). Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis 42*(03), 621–656.

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

Kirkby, J. L., S. Mitra, and D. Nguyen (2020). An analysis of dollar cost averaging and market timing investment strategies. *European Journal of Operational Research 286*(3), 1168–1186.

Klebnikov, S. (2022, May). Warren Buffett's $51 Billion Stock Market Shopping Spree: Here's What He's Buying. https://www.forbes.com/sites/sergeiklebnikov/2022/05/17/warren-buffetts-51-billion-stock-market-shopping-spree-heres-what-hes-buying/?sh=65e3d9df147c.

Ledoit, O. and M. Wolf (2003). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of empirical finance 10*(5), 603–621.

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

Li, J. and C. Wu (2006). Daily return volatility, bid-ask spreads, and information flow: Analyzing the information content of volume. *The Journal of Business 79*(5), 2697–2739.

Loos, B., A. Previtero, S. Scheurle, and A. Hackethal (2020). Robo-advisers and investor behavior. Technical report, Working paper, University of Technology of Sidney.

Malkiel, B. G. (1999). *A random walk down Wall Street: including a life-cycle guide to personal investing*. WW Norton & Company.

Meir, S. (1995). A behavioral framework for dollar-cost averaging. *Journal of Portfolio Management 22*(1), 70–78.

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

Michaud, R. O. (1989). The markowitz optimization enigma: is' optimized'optimal? *Financial Analysts Journal 45*(1), 31–42.

Nam, J.-s. and B. Branch (1994). Tactical asset allocation: can it work? *Journal of Financial Research 17*(4), 465–479.

Nicolosi, G., L. Peng, and N. Zhu (2009). Do individual investors learn from their trading experience? *Journal of Financial Markets 12*(2), 317–336.

Sharpe, W. F. (2010). Adaptive asset allocation policies. *Financial Analysts Journal 66*(3), 45–59.

Sheedy, E., R. Trevor, and J. Wood (1999). Asset-allocation decisions when risk is changing. *Journal of Financial Research 22*(3), 301–315.

Simaan, M. and Y. Simaan (2019). Rational explanation for rule-of-thumb practices in asset allocation. *Quantitative Finance 19*(12), 2095–2109.

Sortino, F. A. and R. Van Der Meer (1991). Downside risk. *Journal of Portfolio Management 17*(4), 27.

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

Figures

Figure 1: **Buy the Dip Popularity Score**
The following plot illustrates the popularity score of searches for Buy the Dip in the 169 months between January 2009 and December 2022. Google defines popularity as the ratio of how many searches have been done for a particular term relative to the total number of searches done on Google. The original data are normalized and presented on a scale from 0-100 where each point on the graph is divided by the highest point, or 100. The green bars below are standardized to reflect a standard deviation of 1. The solid black line denotes the cumulative return on the SPDR S&P 500 ETF Trust (SPY) during the same period. The gray bars denote four different periods. The first two correspond to major market selloffs that took place at the beginning and end of 2018. The third one corresponds to the COVID-19 market selloff, whereas the last denotes the beginning of the bear market that took place in June 2022.

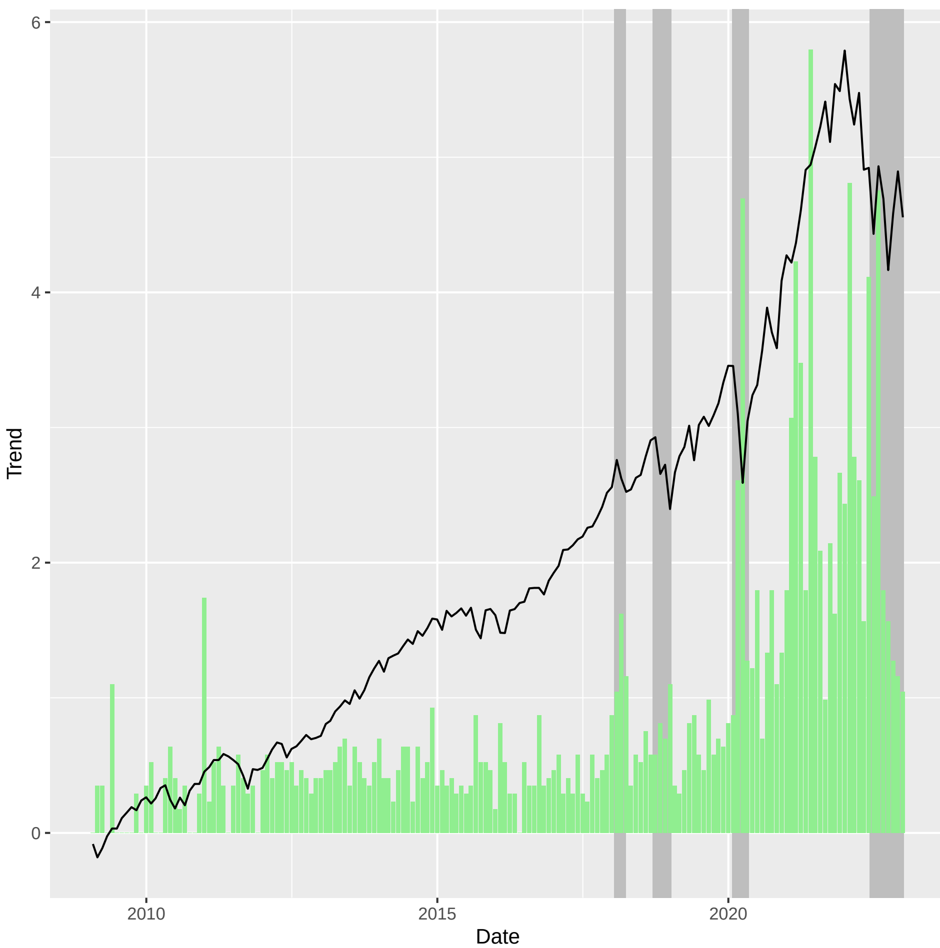

Figure 2: **SPY Cumulative Return**

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 (3.1). Grey bars denote recession periods identified with respect to the National Bureau of Economic Research (NBER).

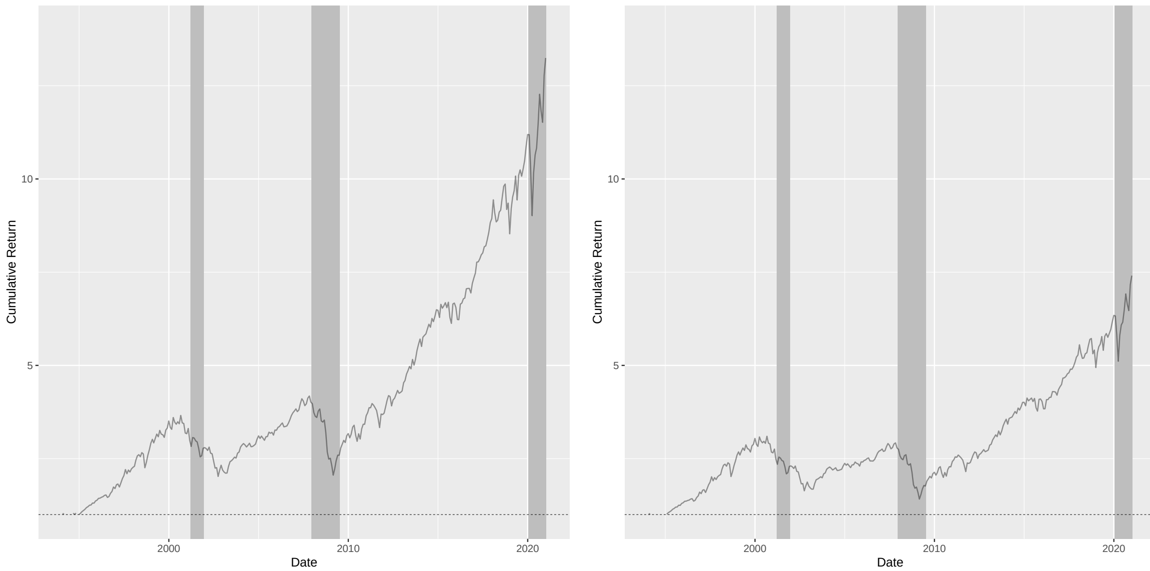

(a) Raw Returns (b) Inflation-adjusted

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).

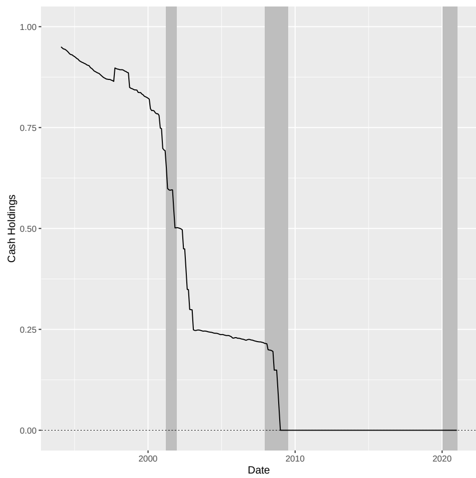

(a) $\tau = -0.05$

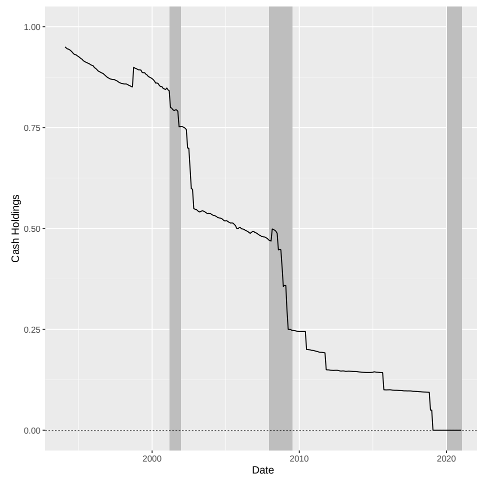

(b) $\tau = -0.06$

Figure 4: **Buy the Dip with Lump Sum - Terminal Wealth Heatmap**
The figure below illustrates terminal wealth for different lump sum BTD investment policies. In particular, it provides a visual perspective on the results in Panel (a) from Table 1. The panels below demonstrate terminal wealth using a heatmap, where red (blue) colors denote high (low) values. The x-axis and y-axis of each figure correspond to the values of k and τ, respectively. The left hand side values of the grid denote a passive strategy that does not follow a BTD policy. Each panel denotes a different testing period. In all cases, the testing period dates are between year y and 2020 (included). For Panels (a) and (b), $y \in \{1994, 2010\}$ denotes periods that begin with bull markets. Similarly, in Panels (c) and (d), $y \in \{2000, 2008\}$ denotes periods that begin with bear markets.

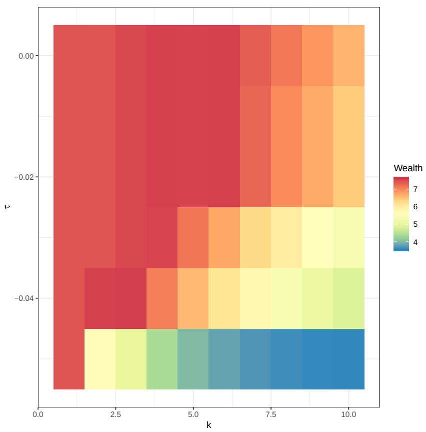

(a) Since 1994

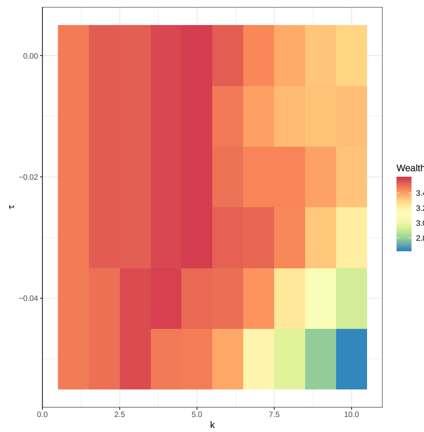

(b) Since 2010

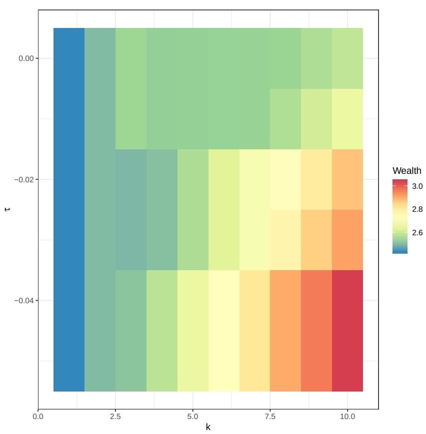

(c) Since 2000

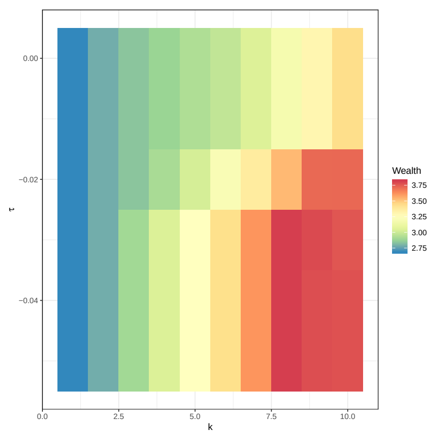

(d) Since 2008

Figure 5: **Performance using Monte Carlo Simulation**

This figure plots the distribution of different performance metrics using a Monte Carlo simulation study. For the simulation, we consider arbitrary values of $\rho \in \{-0.5, -0.4, \cdots, 0, \cdots, 0.4, 0.5\}$ an AR(1) model where ρ denotes the autoregressive parameter. For instance, $\rho = 0$ denotes a random walk, whereas $\rho < 0$ ($\rho > 0$) correspond to reversal (momentum). For a given value of ρ, the simulation conducts the following steps. First, calibrate the parameters such that the unconditional mean return and volatility are consistent across all specifications. Second, generate five years of daily returns using the data process. Third, implement the same BTD experiment from Section 3.3.2 using the simulated data and evaluate performance. In all cases, set $\tau = -0.01$, i.e., allocate one unit of cash to the risky asset each time the daily return drops below -1%. This is repeated 1000 times for each ρ value, resulting roughly 11,000 performance observations. Panels (a), (b), (c), and (d) correspond to the difference in performance in terms volatility, Sortino ratio (Sortino and Van Der Meer, 1991), Value-at-Risk (VaR), and terminal wealth, respectively. Each panel covers the distribution of the performance differential. The distributions are ranked from top to bottom in a descending order with respect to ρ. The middle distribution denotes the case of a random walk, i.e., $\rho = 0$. For each distribution, the plot is trimmed at the 1% and 99% to mitigate extreme outliers. In all cases, the vertical dashed line denotes the zero level.

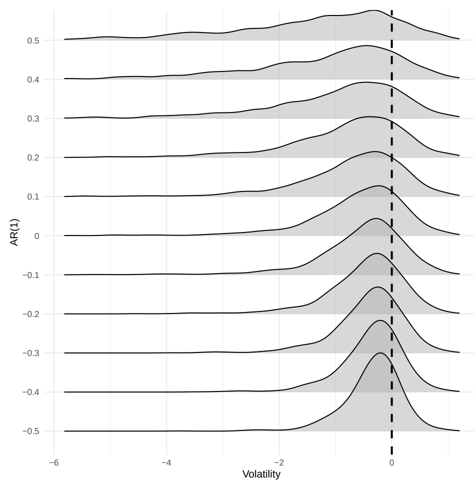

(a) Volatility

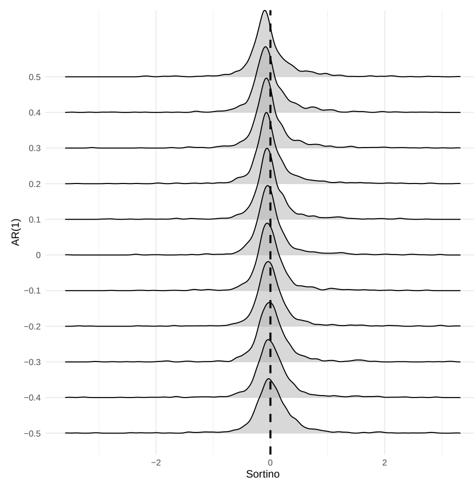

(b) Sortino

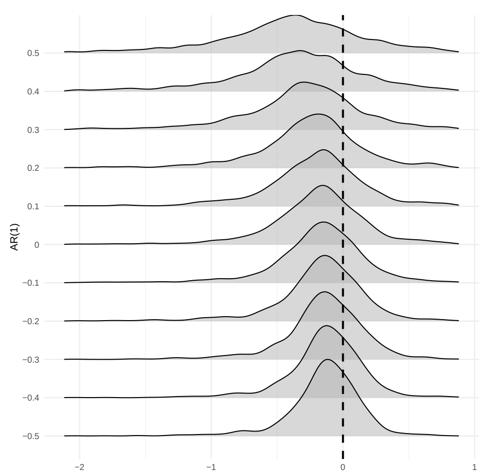

(c) VaR

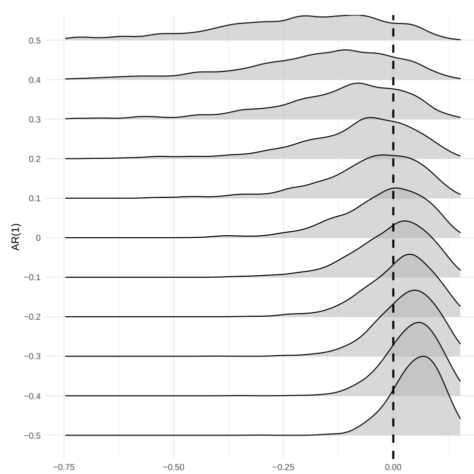

(d) Wealth

Figure 6: **AR(1) Over Time**
This figure illustrates the first order serial correlation, i.e., AR(1) coefficient, using SPY ETF daily returns. 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 (3.1). Grey bars denote recession periods identified with respect to the National Bureau of Economic Research (NBER).

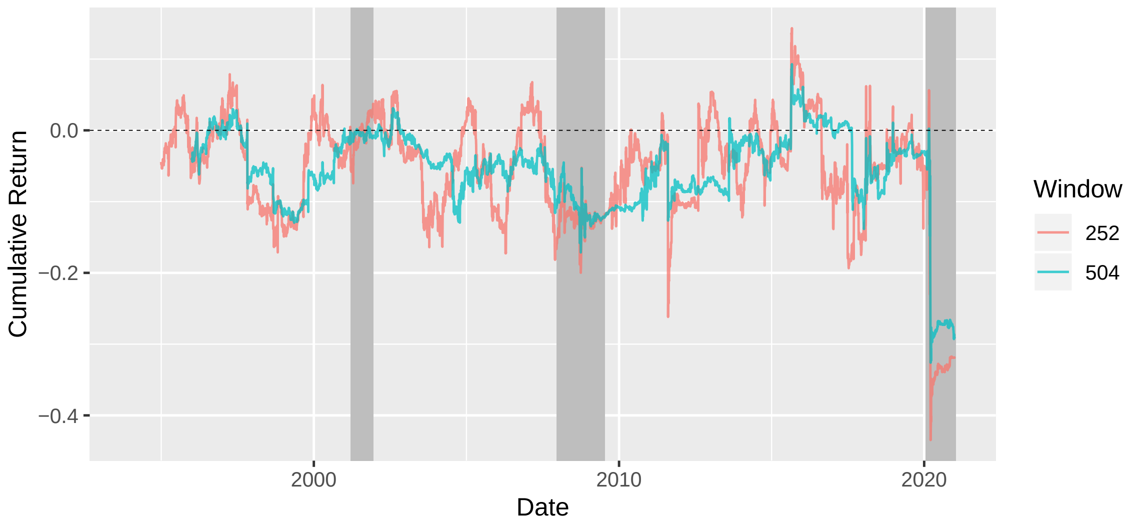

Figure 7: **AR(1) Fixed Effects**
This figure demonstrates the fixed effects (FE) of the autoregressive parameter ρ based on the panel regression summarized in Table 6. Panel (a) denotes the panel regression where the dependent variable is given by the mean return differential, whereas Panel (b) corresponds to the case where it is given by the wealth differential. The x-axis denotes the value of ρ, which ranges between -0.5 and 0.5 using an increment of size 0.1. The y-axis denotes the fixed effect of the panel regression, i.e., the effect of the parameter ρ on the results.

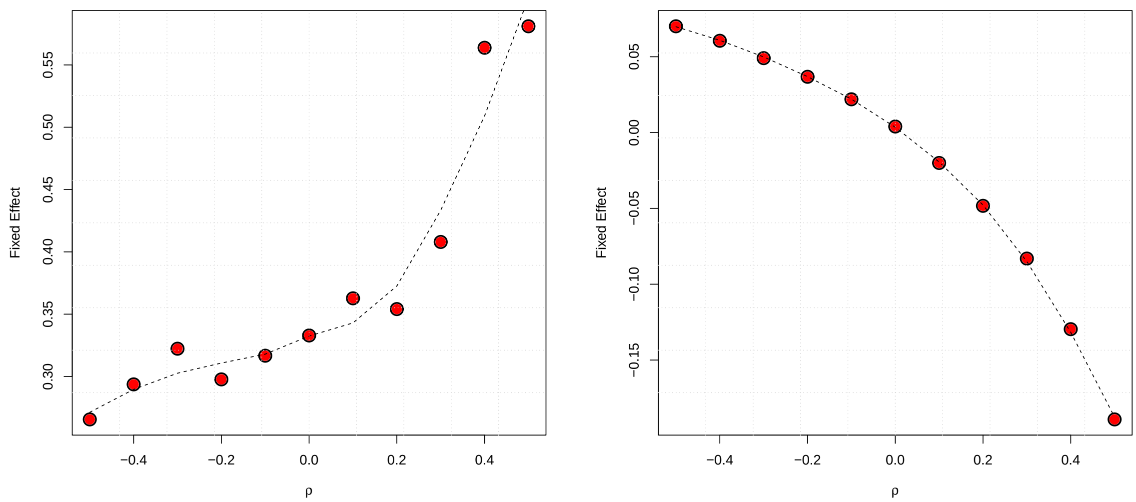

(a) Mean Return (b) Terminal Wealth

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 are 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 date are 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 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 windows, 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 are 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 (0) whether the start of the backtesting period corresponds to a bull (bear) period. The Passive column denotes the allocation policy that invests funds in 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

Table 4: **Information Flow Surrounding Dip Periods**

	Full Sample			*Sub-Sample*		
	Before	After	Difference	Before	After	Difference
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	-0.016^{***}	-0.017^{***}	-0.001	-0.032^{***}	-0.014^{***}	0.018^{***}
	(0.001)	(0.001)	(0.001)	(0.003)	(0.003)	(0.005)
Observations	6,389	6,389	6,389	306	306	306

Note: *p<0.1; **p<0.05; ***p<0.01

Table 5: **Monte Carlo Simulation Summary Statistics**

This table reports the median of different performance metrics using a Monte Carlo simulation study. The simulation considers arbitrary values of $\rho \in \{-0.5, -0.4, \cdots, 0, \cdots, 0.4, 0.5\}$ for an AR(1) model where ρ denotes the autoregressive parameter. For instance, $rho = 0$ denotes a random walk, whereas $\rho < 0$ ($\rho > 0$) corresponds to reversal (momentum). For a given value of ρ, the simulation conducts the following steps. First, calibrate the parameters such that the unconditional mean return and volatility are consistent across all specifications. Second, generate five years of daily returns using the data process. Third, implement the same BTD experiment from Section 3.3.2 using the simulated data and evaluate performance. In all cases, set $\tau = -0.01$, i.e., allocate one unit of cash to the risky asset each time the daily return drops below -1%. This is repeated 1000 times for each ρ value, resulting in roughly 11,000 performance observations. Columns 2, 3, 4, and 5 correspond to the difference in performance in terms volatility, Sortino ratio (Sortino and Van Der Meer, 1991), Value-at-Risk (VaR), and terminal wealth, respectively. The rows are ranked from top to bottom in a descending order with respect to ρ.

ρ	Volatility	Sortino	VaR	wealth
0.50	-1.12	-0.08	-0.38	-0.21
0.40	-0.85	-0.06	-0.32	-0.15
0.30	-0.73	-0.05	-0.27	-0.10
0.20	-0.60	-0.05	-0.22	-0.06
0.10	-0.53	-0.03	-0.18	-0.04
0.00	-0.45	-0.03	-0.17	-0.01
-0.10	-0.38	-0.02	-0.15	0.01
-0.20	-0.35	-0.01	-0.14	0.03
-0.30	-0.32	0.01	-0.13	0.04
-0.40	-0.29	0.01	-0.11	0.05
-0.50	-0.27	0.01	-0.11	0.06

Table 6: **Monte Carlo Simulation Regression Results**

This table reports regression results using simulated data. In total, there 10,963 observations capturing the difference in performance between BTD and passive strategies. The simulation considers arbitrary values of $\rho \in \{-0.5, -0.4, \cdots, 0, \cdots, 0.4, 0.5\}$ for an AR(1) model where ρ denotes the autoregressive parameter. For instance, $\rho = 0$ denotes a random walk, whereas $\rho < 0$ ($\rho > 0$) corresponds to reversal (momentum). For a given value of ρ, the simulation conducts the following steps. First, calibrate the parameters such that the unconditional mean return and volatility are consistent across all specifications. Second, generate five years of daily returns using the data process. Third, implement the same BTD experiment from Section 3.3.2 using the simulated data and evaluate performance. In all cases, set $\tau = -0.01$, i.e., allocate one unit of cash to the risky asset each time the daily return drops below -1%. This is repeated 1000 times for each ρ value, resulting in roughly 11,000 performance observations. Columns (1) and (2) correspond to ordinary least squares (OLS) regressions, where the dependent variable is the difference in mean return and terminal wealth, respectively. Columns (3) and (4) replicate (1) and (2), however, using a panel regression that includes fixed effects for ρ. In all cases, the regressor (independent variable) is the volatility differential.

	Dependent variable:			
	Mean Return	Terminal Wealth	Mean Return	Terminal Wealth
	OLS	*OLS*	*Panel OLS*	*Panel OLS*
	(1)	(2)	(3)	(4)
Volatility	0.683***	0.069***	0.707***	0.049***
	(0.024)	(0.001)	(0.025)	(0.001)
Observations	10,963	10,963	10,963	10,963
Adjusted R^2	0.068	0.248	0.065	0.165
Note:				*p<0.1; **p<0.05; ***p<0.01