The Value of Economic Regularization for Stock Return Predictability*

Yoontae Jeon† Laleh Samarbakhsh‡

Abstract

We propose a new approach on using stock market return predictors to maximize investor's utility gains. In contrast to the conventional OLS approach using statistical loss function, our approach proposes to use economically motivated loss function. Our approach is computationally easy to implement and delivers superior out-of-sample utility gain measured by the certainty equivalent. It further demonstrates that the advantage becomes larger when an investor considers higher moments of portfolio returns and has a larger degree of risk aversion coefficient. Our result points towards the importance of aligning the loss function with the out-of-sample evaluation metric when using return predictor variables.

JEL Classification:

Keywords: return predictability; regularization; certainty equivalent

*We thank Giovanna Nicodano and participants at the AFFI 2021 conference for helpful comments and suggestions. We also thank Amit Goyal for making the data available.

†Corresponding author, DeGroote School of Business, McMaster University, Contact: jeony5@mcmaster.ca

‡Ted Rogers School of Management, Ryerson University, 350 Victoria Street, Toronto, Ontario, Canada, M5B 2K3, Contact: lsamarbakhsh@ryerson.ca

1 Introduction

The literature on stock market return predictability has one of the longest history in finance. Earlier studies have identified many economically motivated predictors that are shown to predict future stock market returns. In contrast, an influential paper of Welch and Goyal (2008) raises a concern that many of the known stock market predictors do not offer significant predictability out-of-sample, when compared to the simple historical average of past equity premiums. Following their work, stream of works came up with a remedy to suggest statistical methods on how to make better use of the stock market return predictors. Campbell and Thompson (2008) propose an economically motivated constraint to place a restriction on the sign of the coefficient and predicted equity premium, which is shown to bring the predictive power of economic predictors in out-of-sample tests. In parallel, Pettenuzzo et al. (2014) propose a constraint on the conditional Sharpe ratio for the market return to further improve the benefits of the stock market return predictors. The essence of these studies lies on the use of economically motivated constraints that overrules the statistical estimation procedure.

Our paper is motivated by these findings and attempts to extend even further by completely abstracting away from the statistical estimation procedure. Although R^2 is the most popular measure of the statistical fit, it may not be economically meaningful for an investor. Consequently, the economic value offered by the return predictors are often measured by the utility gain, or certainty equivalent, of an investor using the predicted equity premium to allocate her assets between risky and risk-free assets. To demonstrate this point, Campbell and Thompson (2008) show that even small amount of out-of-sample statistical fit measured by R^2 can lead to substantial improvement in out-of-sample certainty equivalent gain. Therefore, it is natural to ask whether an investor should completely neglect the statistical efficiency and focus on utility gains. This is the main intuition of the proposed methodology in this paper.

The traditional OLS approach predicts the next-period equity premium to be a linear function of the predictor variable (Equity Premium$_{t+1} = \alpha + \beta X_t + \epsilon_{t+1}$). We take this functional relationship as granted and propose the estimation of the associated coefficients to be carried out by maximizing in-sample certainty equivalent gain of an investor, instead of minimizing the sum of squared errors as in an OLS case.

Therefore, our approach is to be understood as a modification to the loss function in the linear predictive relationship. Note that there exists multiple different ways to improve OLS model in the statistics literature when OLS is not feasible. These methods, which are often called regularization or broadly referred to as a machine learning algorithm, applies a tweak to the original loss function to make the optimization problem tractable. Our approach is in parallel to such statistical regularization approaches. We focus on tweaking the loss function with the particular focus on maximizing an economic gain of the investor. Therefore, we refer to our method as *economic regularization* in this paper.

We empirically test our approach against OLS method using the 14 return predictors studied in Welch and Goyal (2008). The results indicate substantial benefits for investors aligning the loss function with the evaluation metric. When tested for out-of-sample period up to 2018, a mean-variance investor with the risk aversion coefficient of 3 obtains an average of 0.27% per month of additional certainty equivalent gain by using our proposed method over OLS. The economic gain becomes larger when an investor uses CRRA utility function, delivering 3.52% extra certainty equivalent gain. Overall, our results suggest that the benefit of using economic regularization is greater when an investor has preference over higher moments and has higher degree of risk aversion coefficient. These results directly stem from the setup of our approach that uniquely considers specific choice of utility function and risk aversion coefficient into the estimation procedure.

The contribution of our paper is on the introduction of a new method to make use of return predictors, rather than finding a new predictor variable. Thus, our paper can be seen as an alternative way of re-visiting the time-series return predictability literature where the predictors are often judged to be good or bad based on their ability to improve statistical in-sample and out-of-sample fits. Consequently, many of the return predictors were found to be not useful based on statistical criteria. Our proposed method suggests that it is worthwhile to revisit them and assess whether some of them can deliver superior economic gain when combined with the economically motivated in-sample estimation of coefficients.

Our economic regularization approach offers several advantages over the existing methodologies. First, as we are maximizing the in-sample utility gains directly, our approach delivers different estimates of the coefficients depending on the specific utility functions used, as well as different coefficient of risk aversion. The previous portfolio allocation approach is a two-step procedure where an investor estimates the coefficients by running an OLS model, then uses the plug-in rule to estimate the optimal portfolio weights. The specific choice of the utility function and risk aversion coefficient only comes in the second step. In contrast, our approach combines these two-steps into a one-step procedure that estimates the coefficients while embedding the choice of utility function and risk aversion coefficient.

Second, our approach is also capable of embedding the time-varying volatility of stock market returns. Again, the time-varying volatility is not considered until the second step where the weights are determined in the traditional approach, but our method unifies the time-varying nature of the volatility into the estimation procedure. This can be seen as an alternative to the WLS-EV method proposed in Johnson (2019) that embeds the time-varying volatility into the OLS regression.

Our approach also differs from the stream of literature using more advanced statistical models or different utility functions to increase the power of using stock market return predictors. For instance, in an influential paper of DeMiguel, Garlappi, and Uppal (2009), they find that using minimum variance, 1/N, or Bayesian methods deliver superior performance over the conventional mean-variance approach in constructing optimal portfolio allocation. We differ from this line of approach that we do not attempt to use different utility functions or statistical models, but rather propose to simply use an alternative parameter estimation method by using economically motivated loss function. As a result, our approach is flexible that can be also applied for using different utility functions or using an alternative statistical models, which we demonstrate in the robustness check.

The closest paper to ours is perhaps Brandt, Santa-Clara, and Valkanov (2009). They propose a parametric portfolio choice rule to directly estimate the loading on the individual stock characteristics by maximizing a utility of the investor's portfolio return. Our approach builds upon their intuition to study a similar framework in the stock market return predictability context. Also, while their study focuses on the cross-section of stock returns we study the time-series property of using economically motivated objective function.

The rest of the paper is organized as follows: we provide a detailed description of our proposed methodology in Section 2. Then, we present the empirical results in Section 3. Section 4 discusses various extensions and robustness checks of the result. We conclude in Section 5.

2 Methodology

The traditional OLS approach to forecast stock market returns using an economic predictor is based on the following predictive regression specification

$$\text{Equity Premium}_{t+1} = \alpha + \beta X_t + \epsilon_{t+1}. \tag{1}$$

In the above specification, the coefficients α and β are estimated by regressing lagged predictor $\{X_t\}_{t=1}^{T-1}$ on the next period's equity premium $\{R_{t+1} - R_t^f\}_{t=1}^{T-1}$, where R_{t+1} is the stock market return between time t and $t+1$ and R_t^f is the risk-free rate prevailing at time t for the same period. In estimating any econometric models, one has to specify the loss function to state the optimization problem to be solved. In the case of OLS model, the corresponding loss function is the sum of squared residuals from fitted relationship. In the form of an equation, estimating the coefficients of OLS model can be written as the solution to the following optimization problem

$$\min_{\alpha,\beta} \sum_{t=1}^{T-1} (R_{t+1} - R_t^f - \alpha - \beta X_t)^2. \tag{2}$$

Although OLS approach is the most intuitive way of studying the linear predictive relationships in econometric sense, previous literature has found questionable empirical results on the out-of-sample (OOS) performance of known economic predictors. Welch and Goyal (2008) provide a comprehensive study to test if any of the economic predictors can beat simple historical average of past equity premiums in the out-of-sample equity premium prediction. Their conclusion is largely pessimistic in the sense that they don't find statistically significant predictors in out-of-sample for those predictors shown to deliver in-sample predictability.

In order to overcome this issue, Campbell and Thompson (2008) propose an economically constraint approach to modify the standard OLS regression. Campbell and Thompson (2008) suggest that by using the following simple adjustment: 1) whenever the estimated sign of β is not consistent with the economic theory, then use the historical average of past equity premium as a next-period prediction instead, and 2) whenever the predicted next-period equity premium is negative, then set it equal to 0 instead.

These two simple "economic" modifications to OLS approach are then shown to greatly improve the out-of-sample performance of economic predictors. On the other hand, Pettenuzzo, Timmermann, and Valkanov (2014) propose an alternative method by requiring conditional annualized Sharpe ratios for the market return to be bounded between 0 and 1. The proposed method of Pettenuzzo, Timmermann, and Valkanov (2014) is then shown to further improve the out-of-sample predictive power on top of Campbell and Thompson (2008), thus further providing assurance on the usefulness of economic constraint.

Note that, in the aforementioned studies, the metric used to evaluate out-of-sample performance is based on regression's R^2 that measures the portion of variability in equity premium explained by the economic predictor. Again, although this is the most appealing metric in the statistical sense, it may not directly translate to the value for a stock market investor who is using the model. To address this issue, several studies shifted their focus to the economic value of using predictors for forecasting stock market returns.[1] Specifically, instead of using R^2 as the measure of usefulness for each predictor, certainty equivalent (CE) measure derived from utility functions were introduced to understand the economic gain of an investor using return predictors. Campbell and Thompson (2008) suggest that even very small out-of-sample R^2 value can lead to substantial economic gains.

Connecting the two observations above, economic constraint and economic evaluation criteria, we propose a new approach for using stock market predictors, which we name *economic regularization* method. Our proposed approach estimates the coefficients α and β, similar to that in Equation 1, by directly maximizing in-sample certainty equivalent, rather than minimizing OLS loss function. In other words, we rely on economic loss function instead of statistical loss function to maximize out-of-sample economic gain.

Prior studies suggest the importance of aligning the loss function used for estimation and evaluation. For example, Engle (1993) points out the importance of choosing loss function when defining a new model, and Christoffersen and Jacobs (2004) show the importance of aligning the loss function with the evaluation metric for the option pricing models. The motivation of our approach comes from the same intuition that if an investor is interested in potential economic gains out-of-sample, then the in-sample estimation should use a loss function that maximizes economic benefit.

In a recent work, Cederburg, Johnson, and O'Doherty (2019) demonstrate that statistically strong predictors can be sometimes economically unimportant if they tend to take extreme values in high-volatility periods, have low persistence, or follow distributions with fat tails. These findings further supports the need of aligning loss functions for the esti-

[1] Cenesizoglu and Timmermann (2012) compare the economic value of multiple return prediction models. See Rapach and Zhou (2013) for comprehensive summary of literature. Neely et al. (2014) use technical indicators in forecasting equity risk premium while Rapach et al. (2016) document short interest as a strong predictor with out-of-sample economic value. Huang and Zhou (2017) and Mei and Nogales (2018) use CRRA utility functions in their predictability models, among others.

mation and evaluation, as it suggests that statistically strong predictors in-sample does not necessarily translate to the superior out-of-sample performance economically.

As an example, consider an investor who has mean-variance (MV) utility function with risk aversion coefficient γ. She faces the portfolio allocation problem where she has to find optimal weight between stock market and risk-free asset. At each time t, it can be easily shown that the optimal weight of her portfolio in the risky asset, the stock market, is given by

$$w_t = \frac{1}{\gamma} \frac{E_t[R_{t+1} - R_t^f]}{\sigma_{t+1|t}^2}, \tag{3}$$

where $E_t[R_{t+1} - R_t^f]$ is the conditional expectation of the next period's equity premium and $\sigma_{t+1|t}^2$ is the conditional forecast of the next period's stock return variance. Investors using OLS model will thus use $E_t[R_{t+1} - R_t^f] = \hat{\alpha} + \hat{\beta} X_t$ when using X_t as the economic predictor variable. The certainty equivalent metric corresponds to the utility gain of an investor who is using this expression to make a portfolio allocation out-of-sample. In contrast, our proposed economic regularization method aims to directly estimate $\hat{\alpha}$ and $\hat{\beta}$ by maximizing in-sample certainty equivalent. Specifically, our method can be summarized as the following counterpart to the Equation (2)

$$\max_{\alpha,\beta} \frac{1}{T-1} \sum_{t=1}^{T-1} (w_t R_{t+1} + (1 - w_t) R_t^f) - \frac{\gamma}{2} \frac{1}{T-2} \sum_{t=1}^{T-1} (w_t R_{t+1} + (1 - w_t) R_t^f - \bar{R})^2, \tag{4}$$

where $\bar{R} = \frac{1}{T-1} \sum_{t=1}^{T-1} w_t R_{t+1} + (1 - w_t) R_t^f$ is the sample mean of the portfolio return. The above expression is simply an in-sample estimate of the unconditional utility for mean-variance investor. Therefore, we are estimating the coefficients α and β such that it maximizes in-sample gain of the investor who is using linear rule ($E_t[R_{t+1} - R_t^f] = \alpha + \beta X_t$) to allocate her optimal weights.

Our approach drastically differs from the OLS approach. Especially, we do not make use of any statistical metric to define the loss function, but rather fully rely on the economic intuition. One way to understand the rationale behind our approach is perhaps by looking at the *statistical regularization* methods gaining popularity in the recent literature, which are more often coined with the term Machine Learning (ML). Various ML techniques are proposed in order to overcome statistical issues with OLS framework. Many of them apply a tweak to the loss function, as in Equation (2), by adding an additional term or changing the metric, so that the estimation can be carried out where OLS approach is not feasible.[2] Our approach is similar to ML techniques in nature, as we also apply a tweak to the loss function to estimate the coefficients. On the other hand, the biggest difference between our approach and ML approach lies in the motivation. While ML techniques are proposed to overcome statistical issues that make OLS method not feasible, our proposed approach tries to maximize economic benefit to the users of the method. Therefore, we name it *economic regularization* instead.

The intuition of our approach is similar to Brandt et al. (2009), who propose parametric portfolio choice rule to construct optimal portfolio in the cross-section of stocks. In their approach, different characteristics of individual stocks are used to construct a parametric portfolio by solving a similar utility maximization problem in-sample. Our approach can be seen as a time-series extension (using equity premium predictors), of their cross-sectional application (using individual stock characteristics).

Lastly, our approach also carries the same spirit as the WLS-EV approach of Johnson (2019). Johnson (2019) demonstrate that using the weighted least squares approach, taking the time-varying variance into consideration, significantly improves the out-of-sample R^2 of the stock return predictive regression. We would like to emphasize that our approach also embeds the time-varying stock market variance implicitly through the functional form of the optimal weights. When constructing the in-sample objective function to be maximized, portfolio weights on the risky asset are determined by the ratio of equity premium forecast and stock market variance. Hence, our objective function automatically takes the time-

[2]Some papers in this direction for finance literature include Rapach, Strauss, and Zhou (2013) who first introduce LASSO and Gu, Kelly, and Xiu (2020) who study various ML methods in stock return predictions, among others.

varying nature of the stock market variance into consideration by varying the weight placed on the risky asset at each period.

Therefore, previous literature suggests rationale behind our proposed methodology and it is now a purely empirical question of whether the proposed approach will perform better than the OLS approach. In the next section, we take our approach to real data and study whether the proposed methodology can deliver superior certainty equivalent gain out-of-sample.

3 Empirical Results

3.1 Data

We use the same set of economic predictors studied in Welch and Goyal (2008).[3] We download the data up to Dec. 2018 from Amit Goyal's website. From their original paper, we use 14 variables for the empirical study: divided price ratio (dp), dividend yield (dy), earnings price ratio (ep), dividend payout ratio (de), stock variance (svar), treasury bill rate (tbl), long term yield (lty), long term rate of return (ltr), term spread (tms), default yield spread (dfy), inflation (infl), log of book to market ratio (bm), cross-sectional premium (csp), and net equity expansion (ntis). Table 1 provides descriptive statistics of the variables during our sample period.

Equity premium is defined in the common way as the difference between the realized return of the value-weighted CRSP firms and the prevailing risk-free rate. All returns are transformed to log-return before taking the difference. Since the focus of paper lies on the usefulness of return predictors for the equity premium, we simply assume the next-period forecast of market variance, $\sigma^2_{t+1|t}$, is given by the past-month stock variance σ^2_t. This

[3]Recent literature has identified several economic variables that exhibit stronger return predictability then Welch and Goyal (2008)'s variables. These include the variance risk premium of Bollerslev et al. (2009), aggregate short interest of Rapach et al. (2016), and aggregate implied volatility spread of Han and Li (2020), among others. Since the focus of our paper is placed on the methodology rather than specific variables, we do not consider them in this paper. However, the proposed methodology is directly applicable to any set of predictors found in the literature.

assumption is made throughout the paper when the conditional variance forecast of the variance is needed to compute the optimal weights of the portfolio.

3.2 Main Results

To measure the additional economic benefits our method can bring, we proceed as follows. Following the literature, we use the first 240 months of observations as the starting period. At each time t, if there are more than 240 observations available, we estimate the coefficients α and β for each of the 14 predictors using Equation (4). The estimation is done on a rolling basis with the expanding window. We assume the conservative estimate for the coefficient of risk aversion γ and set it equal to 3.[4]

Then, we compute the out-of-sample certainty equivalent gain of MV investor using the same expression as in Equation (4) except using the out-of-sample observations. We denote this out-of-sample certainty equivalent using our method as CE(CE). In parallel, we also compute two other certainty equivalent gains. First, we denote CE(OLS) for the CE gain of an investor using OLS method to allocate her assets. Second, we denote CE(Hist) for the CE gain of an investor using the simple historical average of past equity premiums to allocate her assets, without using economic predictors.

Two questions need to be examined empirically. First, do CE(CE) and CE(OLS) outperform CE(Hist)? That is, does use of economic predictors deliver superior out-of-sample economic gains for a mean-variance investor? Second, does our proposed methodology outperform the traditional OLS method? To see this clearly, we introduce the following notation

$$\Delta CE(OLS) = CE(OLS) - CE(\text{Hist}) \tag{5}$$
$$\Delta CE(CE) = CE(CE) - CE(\text{Hist}).$$

Table 2 reports the above values for the 14 economic predictors, as well as their differences

[4] In Section 4, we demonstrate that increasing the coefficient of risk aversion actually delivers more favourable result to our approach.

defined by $\Delta CE(CE) - \Delta CE(OLS)$. First, the average out-of-sample extra CE gain of OLS method over historical average is negative, being -0.23%. This means that the investors using the historical average would have been better off, even in terms of the economic gain, than using OLS method with the predictors. Out of 14 predictors, 8 of them exhibit negative extra CE gains and 6 of them show positive values. The range of extra CE gain is moderate, from the lowest of -1.22% using csp (cross-sectional premium) to the highest of 0.65% using ntis (net equity expansion). On the other hand, out-of-sample extra CE gain of the our method over historical average is slightly positive, average being 0.04%, thus delivering definitely better performance than the OLS method, and potentially better performance than the historical average. Interestingly, the variation in the extra CE gain becomes much higher. The worst performing predictor is now lty (long term yield) where it has -3.76% of $\Delta CE(CE)$ value. The best performing predictor, on the other hand, is tms (term spread) that delivers 2.31% of extra CE gain. Out of 14 predictors, 6 of them show negative extra CE gain while 8 of them are positive using the proposed method.

When comparing OLS method to our method, the average difference is 0.27% where only 4 out of 14 predictors show worse performance than the OLS method. The difference is largely driven by the observation coming from lty (long term yield) that has a rather extreme value of -3.56%, but other predictors show moderate levels of improvement over OLS method. Overall, the result of Table 2 suggests the potential advantage for a mean-variance investor by using our method over OLS method. It is worthwhile to note that our method tends to deliver more extreme outcomes than the OLS method, for both gains and losses, that may be appealing to some investors. It looks like as if our proposed approach is working to leverage the economic gains while providing higher average performance.

To understand this divergence between the two approaches better, we report the time-series average of the estimated coefficients α and β in Table 3. Our main interest is the magnitude of estimated loading on the predictor variable, $\bar{\beta}$. Clearly, we see a much larger average for the estimated coefficient β using our method, labelled $\bar{\beta}$ CE, compared to the

OLS method, labeled $\bar{\beta}$ OLS. The differences are sometimes quite extreme. For example, when using svar (stock variance) as the predictor, the average OLS loading is only 0.270 while the proposed method requires the average loading of 3.950. In other words, the in-sample data calls for greater importance to be placed on the economic predictors when maximizing economic gain than statistical gain. The result is somewhat expected since we are aligning the loss function to maximize the economic gain as well, but the enormous differences in the estimate coefficients are somewhat surprising in magnitude.

Figure 1 plots the rolling estimates of β coefficients for the OLS approach and our approach (labelled as CE Approach). Noticeable difference between two approaches immediately arise. The β coefficient estimated using CE approach shows a greater degree of variation across time while OLS approach produces a stable pattern. In fact, the estimate β from CE approach starts from significantly negative value, in contrary to what economic theory suggests, at the early period, but it reverts back to positive with higher loading than in the case of OLS. It suggests that CE approach likely adjusts to the fluctuations in the business cycle faster than the OLS approach. As the OLS approach places an equal weight to all observations equally while the CE approach is able to directly apply weighting of each period through the embedded stock market variance component, CE is able to capture the time-varying nature of the predictive relationship more effectively. Also, the fact that CE approach starts with negative β at the early period suggests further room for improvement by applying economic constraint of Campbell and Thompson (2008), which we later verify in Section 4.

Overall, our first set of empirical results point towards the benefit of aligning in-sample loss function and out-of-sample evaluation metric. In the next Section, we provide further extensions and robustness tests to validate our methodology.

4 Robustness

In this section, we conduct various robustness tests by using a different utility function, applying the economic constraint of Campbell and Thompson (2008), and using various coefficients of risk aversion. All robustness results largely support, indeed even stronger, the usefulness of economic regularization in maximizing out-of-sample CE gains for an investor.

4.1 CRRA Utility

Although mean-variance utility is one of the most popular choices for an optimal asset allocation problem, due to its analytical tractability, it suffers a drawback of not being able to capture risk appetite on higher moments. In general, investors not only care about mean and variance, but rather are concerned with the entire distribution of their portfolio returns. One parsimonious way to address this issue is to use the CRRA utility function defined by the following period-by-period utility

$$U(r_t) = \frac{(1+r_t)^{1-\gamma}}{1-\gamma}. \tag{6}$$

Under the CRRA utility framework, our optimization problem in Equation (4) can be stated as

$$\max_{\alpha,\beta} \sum_{t=1}^{T-1} U(w_t R_{t+1} + (1-w_t)R_t^f), \tag{7}$$

where w_t has the same functional form as in Equation (3).[5] Following Faias and Santa-Clara (2017), we define the out-of-sample certainty equivalent by the following equation

$$CE = [(1-\gamma)\bar{U}]^{1/(1-\gamma)} - 1, \tag{8}$$

[5] Campbell and Viceira (2002) derive the same functional form under the log-normally distributed returns and linear approximations, which we follow.

where $\bar{U}$ is the average CRRA utility of the out-of-sample portfolio returns. We use the same risk aversion coefficient $\gamma = 3$ as in the mean-variance case.

Table 4 reports the results in a similar format to Table 2. Surprisingly, the difference is significantly more noticeable in the CRRA case. The OLS model performs miserably by having an average extra CE gain of -1.94% (or loss of 1.94%) over the simple historical average model. Note that the large extra CE loss is mainly driven by extreme observations, however. Total of 9 out of 14 predictors still show improvement over the historical average. Nevertheless, the performance of our method is quite promising in comparison. Not only is the average extra CE gain is large being 1.58%, but 13 out of 14 predictors show substantially large improvement over the historical average. Moreover, the improvement is quite uniform across all predictors, rather than disperse like the mean-variance case. As a result, a CRRA investor using our method benefits by an average of 3.52% across 14 predictors over an investor using OLS model. Our results highlight the possibility that the true value of using the economic regularization method may lie in the power of optimizing over the higher moments of portfolio return. While the mean-variance utility case still delivers significant differences between the two methods, the CRRA case thus further strengthens the importance of aligning the loss function with the evaluation metric.

Figure 2 plots the time-series of period-by-period CRRA utility of out-of-sample optimal portfolio returns. It is clear from the figure why investors would prefer the CE approach over the OLS approach. The OLS approach delivers extremely volatile time-series of utility gains while the CE approach delivers more stable utilities throughout the out-of-sample period. Although the OLS approach sometimes delivers higher utilities during some periods, their losses relative to the CE approach during bad times are significantly larger. As a result, when CRRA utility is used, investors show greater willingness to use the CE maximization approach in comparison to the OLS approach. Moreover, our approach delivers significantly higher CE gain over the simple historical average approach as well.

4.2 Campbell and Thompson Adjustment

Since our approach is based on maximizing economic gain, it is natural to also consider if our approach is merely capturing the existing economic constraint methods. To verify this, we apply the economic constraint of Campbell and Thompson (2008) (CT) to see whether it provides further improvements. Table 5 reports the result for a mean-variance investor with the risk aversion coefficient of $\gamma = 3$. Comparing the result with Table 2, without the CT adjustment, we see some additional improvements. The average extra CE gain for the OLS method slightly improves from -0.23% to -0.15%, and the cross-sectional dispersion across predictors are reduced. For the case of our method, the improvement is even greater from 0.04% to 0.75%. Also, 12 out of 14 predictors now show positive extra CE gain over the historical average while 8 out of 14 predictors still suffer from negative extra CE gain in the OLS model.

Once again, the improvements mainly come from stabilizing the dispersion across different predictors. Hence, the results suggest that while our proposed methodology focuses on maximizing economic gain, having an additional economic constraint helps to reduce the possible extreme outlier observations. Intuitively, as discussed earlier, since our approach delivers much larger loading coefficient β, having an economic constraint provides a downside protection during the period that the sign of the prediction is wrong.

4.3 Different Risk Aversion Coefficient

Lastly, to ensure our results are not driven by the specific choice of risk aversion parameter $\gamma = 3$, we consider a higher order of risk aversion $\gamma = 5$ and $\gamma = 10$. In principle, our method should work even better with the higher order of risk aversion coefficient because our method directly incorporates the parameter γ in the estimation process. In contrast, OLS method is based on a purely statistical framework, leaving no room for any specific risk aversion coefficient to play a role. Thus, increasing γ, or making the impact of investor's risk appetite more important, should deliver larger differences between our method and the

OLS method.

Table 6 shows the result and indeed it is the case that our method outperforms the OLS model by a large magnitude. We report the extra CE gain our method delivers over the OLS model for two cases of $\gamma = 5$ and $\gamma = 10$. For $\gamma = 5$, the average extra CE gain of our method over OLS model is 1.72% where 11 out of 14 predictors show positive gain. When γ is set equal to 10, the extra CE gain is magnified to average of 4.66% with 12 out of 14 predictors outperforming OLS model. Hence, we conclude that our proposed method is going to be even more useful for investors with higher risk-aversion.

Next, we test whether our proposed method's advantage is dependant on the economic regimes. We split the sample into sub-periods of economic expansion and contraction. Table 7 compares the mean values of the 14 economic predictors from 1946 – 2018 during expansion periods. According to the National Bureau of Economic Research, an expansion period is deemed a "normal state of the economy"[6]. Both empirical models (CE and OLS) exhibit positive differences and all statistically significant except cross-sectional premium (0.06%). The average mean difference during the expansion period of all economic predictors is 3.57%. The largest spread in mean differences is shown in the long-term return rate and term spread (6.02%), statistically significant. During expansion periods, the certainty equivalent measure allows investors to understand economic gain through return predictors. Also, we find that economic predictors deliver superior gains for a mean-variance investor during expansion periods.

Table 8 compares the mean values of the 14 economic predictors from 1946 – 2018 during contraction periods. From literature, a contraction period occurs between the peak and trough of an economy. The mean difference between models during the contraction period is marginal at 1.11% (3.61% vs. 3.57%). During this economic life cycle, the log of book-to-market ratio and cross-sectional premium are not statistically significant in differences. Similarly, during the expansion stage, the long-term rate of return exhibits a spread of

[6]https://www.nber.org/research/business-cycle-dating

11.90%. We see 13 out of 14 predictors show positive extra CE while only 1 out of 14 predictors still suffer from negative CE gain in the OLS model.

The findings of sub-period analysis further support the usefulness of our model to the OLS model regardless of economic climate.

5 Conclusion

This paper proposes a new method to use stock market return predictors, that can enhance out-of-sample utility gains of an investor. Instead of relying on the statistical OLS method, the proposed method seeks to maximize in-sample certainty equivalent of an investor who is using the linear function of return predictor to allocate her assets between risky and risk-free asset. The empirical findings support the idea and shows significant out-of-sample improvements in certainty equivalent gains using the proposed method over OLS or simple historical average. The finding is even more strongly supported when using CRRA utility, economic constraint, or higher order of risk aversion.

Our findings contribute to the literature on the return predictability by proposing a new way of thinking how the economic predictors can be used. The approach is general enough that it can be extended to different classes of utility functions as well as non-linear functional forms. In summary, our results suggest a potential future development on the economically motivated loss functions in a broader context.

References

Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. Review of Financial Studies 22, 4463–4492.

Brandt, M. W., Santa-Clara, P., Valkanov, R., 2009. Parametric portfolio policies: Exploring characteristics in the cross-section of equity returns. Review of Financial Studies 22, 3411–3447.

Campbell, J. Y., Thompson, S. B., 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? Review of Financial Studies 21, 1509–1531.

Campbell, J. Y., Viceira, L. M., 2002. Strategic asset allocation: Portfolio choice for long-term investors. Oxford University Press .

Cederburg, S., Johnson, T. L., O'Doherty, M. S., 2019. On the economic significance of stock return predictability. Working Paper .

Cenesizoglu, T., Timmermann, A., 2012. Do return prediction models add economic value? Journal of Banking & Finance 36, 2974–2987.

Christoffersen, P., Jacobs, K., 2004. The importance of the loss function in option valuation. Journal of Financial Economics 72, 291–318.

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.

Engle, R. F., 1993. A comment on hendry and clements on the limitations of comparing mean square forecast errors. Journal of Forecasting 12, 642–644.

Faias, J. A., Santa-Clara, P., 2017. Optimal option portfolio strategies: Deepening the puzzle of index option mispricing. Journal of Financial and Quantitative Analysis 52, 277–303.

Gu, S., Kelly, B., Xiu, D., 2020. Empirical asset pricing via machine learning. Review of Financial Studies .

Han, B., Li, G., 2020. Aggregate implied volatility spread and stock market returns. Management Science .

Huang, D., Zhou, G., 2017. Upper bounds on return predictability. Journal of Financial and Quantitative Analysis 52, 401–425.

Johnson, T. L., 2019. A fresh look at return predictability using a more efficient estimator. Review of Asset Pricing Studies 9, 1–46.

Mei, X., Nogales, F. J., 2018. Portfolio selection with proportional transaction costs and predictability. Journal of Banking & Finance 94, 131–151.

Neely, C. J., Rapach, D. E., Tu, J., Zhou, G., 2014. Forecasting the equity risk premium: the role of technical indicators. Management science 60, 1772–1791.

Pettenuzzo, D., Timmermann, A., Valkanov, R., 2014. Forecasting stock returns under economic constraints. Journal of Financial Economics 114, 517–553.

Rapach, D., Zhou, G., 2013. Forecasting stock returns. Handbook of Economic Forecasting 2, 328–383.

Rapach, D. E., Ringgenberg, M. C., Zhou, G., 2016. Short interest and aggregate stock returns. Journal of Financial Economics 121, 46–65.

Rapach, D. E., Strauss, J. K., Zhou, G., 2013. International stock return predictability: What is the role of the united states? The Journal of Finance 68, 1633–1662.

Welch, I., Goyal, A., 2008. A comprehensive look at the empirical performance of equity premium prediction. Review of Financial Studies 21, 1455–1508.

Table 1: Descriptive Statistics of 14 Predictor Variables

	Period	Mean	Median	Std. Dev.	Skewness
dp	1946–2018	−3.223	−3.153	0.439	−0.651
dy	1946–2018	−3.220	−3.149	0.436	−0.679
ep	1946–2018	−2.676	−2.688	0.377	−0.644
de	1946–2018	−0.547	−0.556	0.315	0.840
svar	1946–2018	0.002	0.001	0.005	6.633
tbl	1946–2018	0.034	0.031	0.030	1.094
lty	1946–2018	0.050	0.042	0.027	1.189
ltr	1946–2018	0.005	0.003	0.024	0.578
tms	1946–2018	0.016	0.017	0.013	−0.151
dfy	1946–2018	0.012	0.009	0.007	2.126
infl	1946–2018	0.003	0.002	0.007	0.812
bm	1946–2018	−0.690	−0.627	0.499	−0.476
csp	1957–2002	0.000	0.000	0.002	0.543
ntis	1947–2018	0.017	0.017	0.026	1.612

This table provides descriptive statistics of the variables during our sample period 1946–2018. The 14 predictor variables as per the empirical study by Welch and Goyal (2008), are: divided price ratio (**dp**), dividend yield (**dy**), earnings price ratio (**ep**), dividend payout ratio (**de**), stock variance (**svar**), treasury bill rate (**tbl**), long term yield (**lty**), long term rate of return (**ltr**), term spread (**tms**), default yield spread (**dfy**), inflation (**infl**), log of book to market ratio (**bm**), cross-sectional premium (**csp**), and net equity expansion (**ntis**).

Table 2: Comparison of Out-of-sample Certainty Equivalent Gains

	Period	Δ CE(OLS)	Δ CE(CE)	Diff
dp	1946–2018	−0.78%	−0.58%	0.20%
dy	1946–2018	−0.98%	−0.57%	0.41%
ep	1946–2018	−0.22%	0.35%	0.57%
de	1946–2018	0.11%	−1.05%	−1.17%
svar	1946–2018	−0.13%	1.07%	1.20%
tbl	1946–2018	0.02%	−0.59%	−0.60%
lty	1946–2018	−0.20%	−3.76%	−3.56%
ltr	1946–2018	0.11%	2.10%	1.98%
tms	1946–2018	0.31%	2.31%	2.00%
dfy	1946–2018	−0.33%	0.21%	0.54%
infl	1946–2018	0.11%	1.47%	1.36%
bm	1946–2018	−0.61%	0.69%	1.30%
csp	1957–2002	−1.22%	−2.43%	−1.21%
ntis	1947–2018	0.65%	1.39%	0.74%
Average:		−0.23%	0.04%	0.27%

This table presents statistics on forecast errors by comparing the out-of-sample CE (certainty equivalent) gains for an investor who has mean-variance (MV) utility function with risk aversion coefficient $\gamma = 3$. The out-of-sample certainty equivalent gain of an MV investor is computed using the same expression as in Equation (4) below, but using the out-of-sample observations instead:

$$\max_{\alpha,\beta} \frac{1}{T-1}\sum_{t=1}^{T-1}(w_t R_{t+1} + (1-w_t)R_t^f) - \frac{\gamma}{2}\frac{1}{T-2}\sum_{t=1}^{T-1}(w_t R_{t+1} + (1-w_t)R_t^f - \bar{R})^2,$$

CE(CE) is the out-of-sample certainty equivalent using our method, compute two other certainty equivalent gains. CE(OLS) denotes the CE gain of an investor using OLS method, and CE(Hist) for the CE gain of an investor using the simple historical average of past equity premium to allocate her assets, without using economic predictors. The deltas are defined as $\Delta CE(OLS) = CE(OLS) - CE(\text{Hist})$ and $\Delta CE(CE) = CE(CE) - CE(\text{Hist})$. The 14 variables from Welch and Goyal (2008)'s empirical study are: divided price ratio (**dp**), dividend yield (**dy**), earnings price ratio (**ep**), dividend payout ratio (**de**), stock variance (**svar**), treasury bill rate (**tbl**), long term yield (**lty**), long term rate of return (**ltr**), term spread (**tms**), default yield spread (**dfy**), inflation (**infl**), log of book to market ratio (**bm**), cross-sectional premium (**csp**), and net equity expansion (**ntis**). The **Diff** column shows noticeable differences between the two models across these variables' CE gains.

Table 3: Average Estimated Coefficients for OLS and CE Approach

	$\bar{\alpha}$ OLS	$\bar{\alpha}$ CE	$\bar{\beta}$ OLS	$\bar{\beta}$ CE
dp	0.059	0.159	0.017	0.043
dy	0.074	0.177	0.022	0.049
ep	0.055	0.123	0.018	0.039
de	0.005	0.025	−0.006	0.003
svar	0.006	0.018	0.270	3.950
tbl	0.011	0.045	−0.138	−0.671
lty	0.016	0.050	−0.257	−0.603
ltr	0.007	0.019	0.013	0.854
tms	0.004	0.011	0.212	0.898
dfy	0.002	0.029	0.447	−0.831
infl	0.008	0.027	−0.494	−3.606
bm	0.016	0.030	0.019	0.018
csp	0.003	0.007	2.492	6.754
ntis	0.013	0.026	−0.236	−0.292

This table reports the time-series average of the coefficients α and β is estimated for an investor who has mean-variance (MV) utility function with risk aversion coefficient $\gamma = 3$. It is notable that the magnitude of the average estimated loadings on the predictor variable, $\bar{\beta}$ is much larger when using our CE method, compared to the OLS method. The 14 variables from Welch and Goyal (2008)'s empirical study are: divided price ratio (**dp**), dividend yield (**dy**), earnings price ratio (**ep**), dividend payout ratio (**de**), stock variance (**svar**), treasury bill rate (**tbl**), long term yield (**lty**), long term rate of return (**ltr**), term spread (**tms**), default yield spread (**dfy**), inflation (**infl**), log of book to market ratio (**bm**), cross-sectional premium (**csp**), and net equity expansion (**ntis**).

Table 4: Comparison of Out-of-sample Certainty Equivalent Gains: CRRA Investor

	Period	Δ CE(OLS)	Δ CE(CE)	Diff
dp	1946–2018	1.84%	1.91%	0.07%
dy	1946–2018	1.34%	1.92%	0.58%
ep	1946–2018	−7.11%	2.04%	9.15%
de	1946–2018	−12.15%	1.94%	14.09%
svar	1946–2018	0.79%	2.11%	1.32%
tbl	1946–2018	0.82%	1.80%	0.98%
lty	1946–2018	0.49%	1.14%	0.65%
ltr	1946–2018	−12.89%	1.46%	14.35%
tms	1946–2018	0.74%	2.21%	1.46%
dfy	1946–2018	1.32%	1.96%	0.64%
infl	1946–2018	−0.41%	2.13%	2.54%
bm	1946–2018	0.58%	2.05%	1.47%
csp	1957–2002	0.79%	2.44%	1.65%
ntis	1947–2018	−3.24%	−2.94%	0.29%
Average:		−1.94%	1.58%	3.52%

This table presents statistics on forecast errors by comparing the out-of-sample CE (certainty equivalent) gains for an investor who has CRRA utility function with risk aversion coefficient $\gamma = 3$, with the out-of-sample OLS performance. The CRRA utility function is defined as $U(r_t) = \frac{(1+r_t)^{1-\gamma}}{1-\gamma}$, and the optimization problem is to maximize $\sum_{t=1}^{T-1} U(w_t R_{t+1} + (1-w_t) R_t^f)$. The out-of-sample certainty equivalent is $CE = [(1-\gamma)\bar{U}]^{1/(1-\gamma)} - 1$ where $\bar{U}$ is the average CRRA utility of the out-of-sample portfolio returns. The 14 variables from Welch and Goyal (2008)'s empirical study are: divided price ratio (**dp**), dividend yield (**dy**), earnings price ratio (**ep**), dividend payout ratio (**de**), stock variance (**svar**), treasury bill rate (**tbl**), long term yield (**lty**), long term rate of return (**ltr**), term spread (**tms**), default yield spread (**dfy**), inflation (**infl**), log of book to market ratio (**bm**), cross-sectional premium (**csp**), and net equity expansion (**ntis**).

Table 5: Comparison of Out-of-sample Certainty Equivalent Gains with CT Adjustment

	Period	Δ CE(OLS)	Δ CE(CE)	Diff
dp	1946–2018	−0.68%	0.10%	0.78%
dy	1946–2018	−0.73%	0.22%	0.95%
ep	1946–2018	−0.20%	0.43%	0.64%
de	1946–2018	0.12%	−0.51%	−0.63%
svar	1946–2018	−0.13%	1.07%	1.20%
tbl	1946–2018	0.01%	1.58%	1.56%
lty	1946–2018	−0.15%	0.65%	0.81%
ltr	1946–2018	0.12%	2.01%	1.89%
tms	1946–2018	0.32%	2.29%	1.98%
dfy	1946–2018	−0.33%	0.21%	0.55%
infl	1946–2018	0.11%	1.58%	1.47%
bm	1946–2018	−0.40%	0.69%	1.09%
csp	1957–2002	−0.87%	−1.20%	−0.33%
ntis	1947–2018	0.65%	1.39%	0.74%
Average:		−0.15%	0.75%	0.91%

This table reports the results of comparing the out-of-sample CE (certainty equivalent) gains after applying the Campbell and Thompson (2008) (CT) economic constraint. The presented forecast errors are for a mean-variance (MV) investor with the risk aversion coefficient of $\gamma = 3$. The 14 variables from Welch and Goyal (2008)'s empirical study are: divided price ratio (**dp**), dividend yield (**dy**), earnings price ratio (**ep**), dividend payout ratio (**de**), stock variance (**svar**), treasury bill rate (**tbl**), long term yield (**lty**), long term rate of return (**ltr**), term spread (**tms**), default yield spread (**dfy**), inflation (**infl**), log of book to market ratio (**bm**), cross-sectional premium (**csp**), and net equity expansion (**ntis**).

Table 6: Comparison of Out-of-sample Certainty Equivalent Gains with Different γ Values

	Period	CE(CE)-CE(OLS)	CE(CE)-CE(OLS)
		$\gamma = 5$	$\gamma = 10$
dp	1946–2018	1.09%	2.91%
dy	1946–2018	1.21%	2.52%
ep	1946–2018	2.23%	5.81%
de	1946–2018	−0.28%	1.14%
svar	1946–2018	3.40%	8.16%
tbl	1946–2018	0.57%	2.61%
lty	1946–2018	−4.58%	−7.90%
ltr	1946–2018	4.93%	11.23%
tms	1946–2018	5.17%	12.00%
dfy	1946–2018	2.09%	5.32%
infl	1946–2018	3.91%	9.42%
bm	1946–2018	3.05%	7.04%
csp	1957–2002	−2.09%	−4.13%
ntis	1947–2018	3.41%	9.07%
Average:		1.72%	4.66%

This table tests our model's robustness by comparing the out-of-sample certainty equivalent (CE) gains for different choices of risk aversion values. Two cases with $\gamma = 5$ and $\gamma = 10$ are presented, showing additional CE advantage over OLS, for investors with higher risk-aversion. The 14 variables from Welch and Goyal (2008)'s empirical study are: divided price ratio (**dp**), dividend yield (**dy**), earnings price ratio (**ep**), dividend payout ratio (**de**), stock variance (**svar**), treasury bill rate (**tbl**), long term yield (**lty**), long term rate of return (**ltr**), term spread (**tms**), default yield spread (**dfy**), inflation (**infl**), log of book to market ratio (**bm**), cross-sectional premium (**csp**), and net equity expansion (**ntis**).

Table 7: Mean Comparison: Economic Expansion Period

	Period	CE	OLS	Diff (CE-OLS)	
dp	1946–2018	2.98%	1.10%	1.88%	**
dy	1946–2018	2.74%	0.86%	1.88%	**
ep	1946–2018	5.11%	2.10%	3.01%	***
de	1946–2018	5.77%	2.67%	3.10%	***
svar	1946–2018	6.47%	2.16%	4.32%	***
tbl	1946–2018	5.82%	2.12%	3.70%	***
lty	1946–2018	3.53%	1.90%	1.63%	*
ltr	1946–2018	8.55%	2.54%	6.02%	***
tms	1946–2018	8.65%	2.62%	6.02%	***
dfy	1946–2018	7.33%	1.83%	5.49%	***
infl	1946–2018	6.81%	2.42%	4.39%	***
bm	1946–2018	6.19%	1.51%	4.69%	***
csp	1946–2018	0.90%	0.84%	0.06%	
ntis	1946–2018	7.21%	3.41%	3.80%	***
Average:		5.58%	2.01%	3.57%	

This table tests the mean differences by comparing the two empirical strategies during the expansion period from 1946 - 2018. The 14 variables from Welch and Goyal (2008)'s empirical study are: divided price ratio (**dp**), dividend yield (**dy**), earnings price ratio (**ep**), dividend payout ratio (**de**), stock variance (**svar**), treasury bill rate (**tbl**), long term yield (**lty**), long term rate of return (**ltr**), term spread (**tms**), default yield spread (**dfy**), inflation (**infl**), log of book to market ratio (**bm**), cross-sectional premium (**csp**), and net equity expansion (**ntis**). Statistical significance of mean difference is indicated at 1%, 5%, and 10% levels.

Table 8: Mean Comparison: Economic Contraction Period

	Period	CE	OLS	Diff (CE-OLS)	
dp	1946–2018	5.45%	1.62%	3.84%	***
dy	1946–2018	6.12%	1.89%	4.23%	***
ep	1946–2018	2.74%	1.09%	1.65%	**
de	1946–2018	2.40%	0.77%	1.63%	*
svar	1946–2018	2.46%	0.76%	1.70%	**
tbl	1946–2018	8.05%	1.65%	6.40%	***
lty	1946–2018	7.05%	1.31%	5.74%	***
ltr	1946–2018	12.92%	1.02%	11.90%	***
tms	1946–2018	6.33%	1.60%	4.73%	***
dfy	1946–2018	3.66%	0.94%	2.71%	***
infl	1946–2018	5.65%	1.21%	4.44%	***
bm	1946–2018	2.61%	1.42%	1.19%	
csp	1946–2018	−2.18%	−0.57%	−1.60%	
ntis	1946–2018	2.92%	0.99%	1.92%	**
Average:		4.73%	1.12%	3.61%	

This table tests the mean differences by comparing the two empirical strategies during the contraction period from 1946 - 2018. The 14 variables from Welch and Goyal (2008)'s empirical study are: divided price ratio (**dp**), dividend yield (**dy**), earnings price ratio (**ep**), dividend payout ratio (**de**), stock variance (**svar**), treasury bill rate (**tbl**), long term yield (**lty**), long term rate of return (**ltr**), term spread (**tms**), default yield spread (**dfy**), inflation (**infl**), log of book to market ratio (**bm**), cross-sectional premium (**csp**), and net equity expansion (**ntis**). Statistical significance of mean difference is indicated at 1%, 5%, and 10% levels.

Figure 1: Time-series of β Coefficient Estimated using D/P as the Predictor Variable

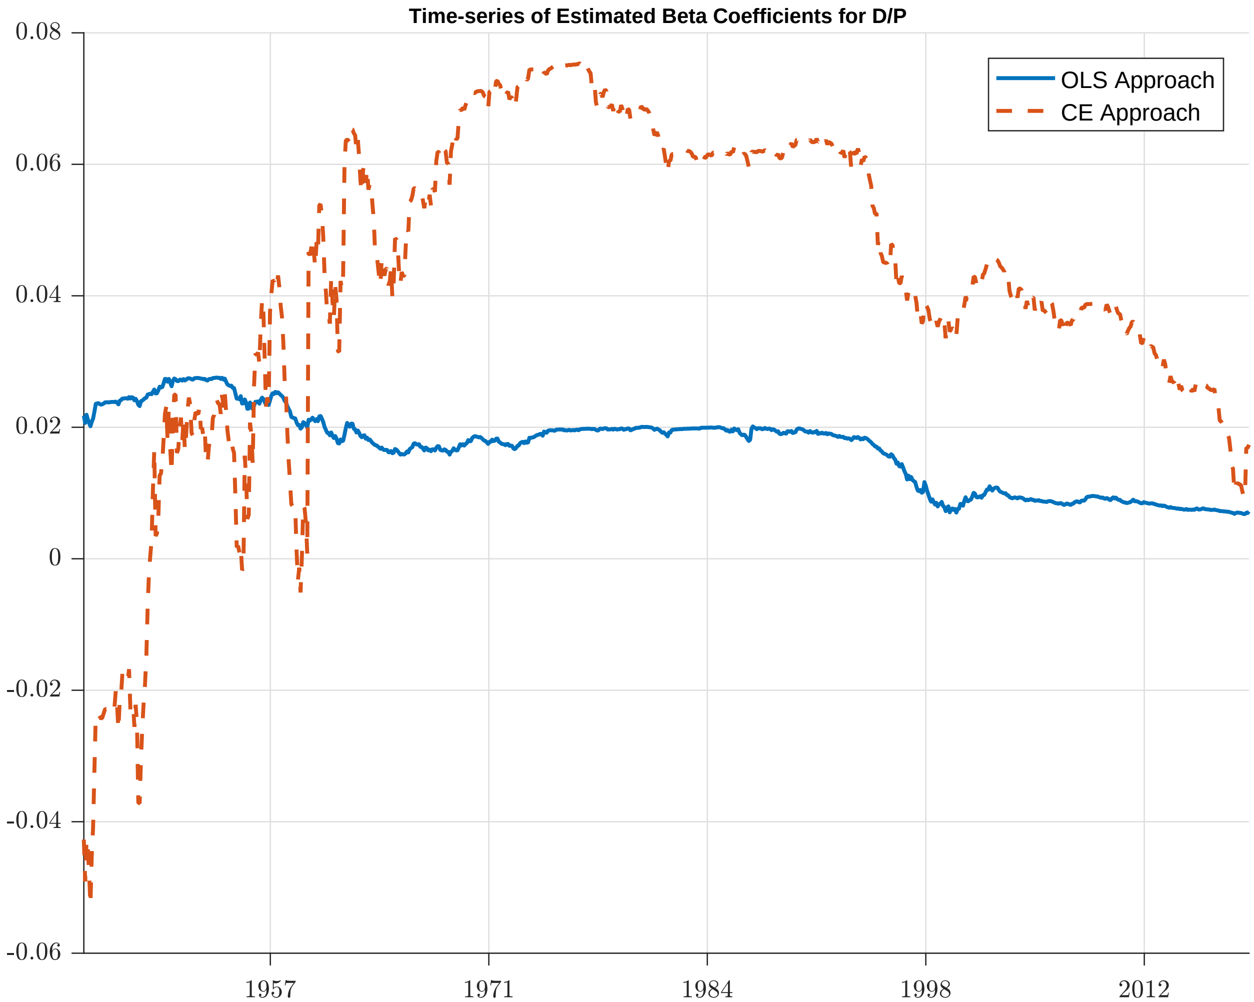

This figure plots the rolling estimates of β coefficients estimated using **dp** (Dividend to Price ratio) as the predictor variable, for the OLS approach (the solid line) vs. our approach (the dashed line) capturing the benefit of aligning in-sample loss function and out-of-sample evaluation metric. The β estimated using CE approach adjusts to the fluctuations in the business cycle (with different weighting for each period from the embedded stock market variance component) faster than the OLS approach (with equal weight to all observations).

Figure 2: Time-series of Out-of-sample CRRA Utility of Optimal Portfolios Constructed using D/P as the Predictor Variable

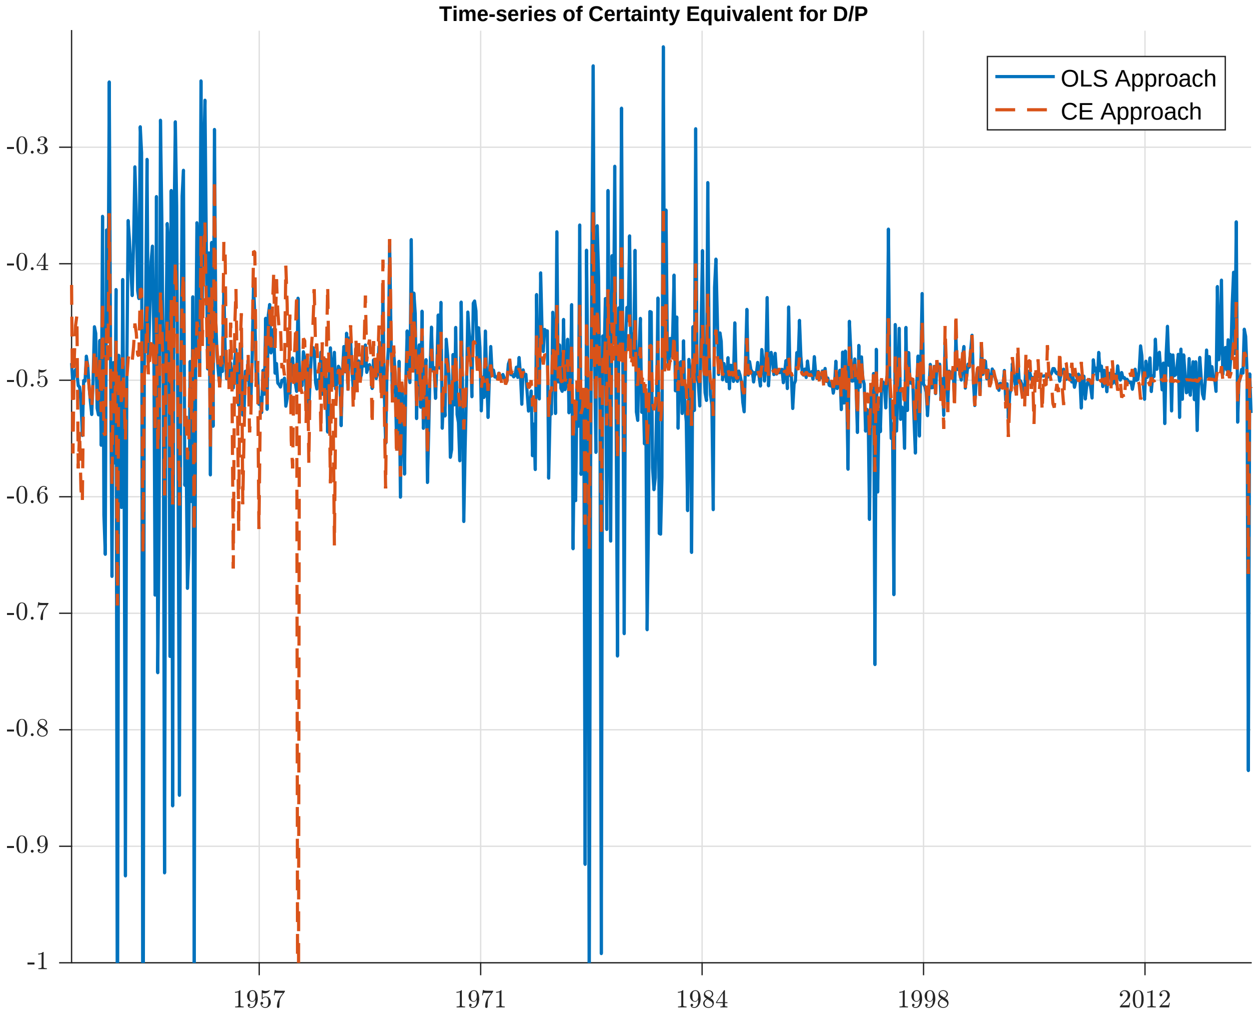

This figure plots the time-series of period-by-period CRRA utility of out-of-sample optimal portfolio returns. The predictor is **dp** (Dividend to Price ratio). The OLS approach (the solid line) delivers extremely volatile time-series of utility gains while CE approach (the dashed line) delivers more stable utilities throughout the out-of-sample period. When CRRA utility is used, investors benefit from using CE maximization approach in comparison to OLS approach. Our CE approach also delivers significantly higher CE gain over simple historical average approach.