2009 American Control Conference
Hyatt Regency Riverfront, St. Louis, MO, USA
June 10-12, 2009

Portfolio Optimization as a Learning Platform for Control Education and Research

Blake Durtschi Mark Skinner Sean Warnick

Abstract— **This paper demonstrates the use of discrete time portfolio optimization as a mechanism for introducing students to key problems in systems theory: control, system identification, model reduction, and verification. Too often students are not introduced to systems theory until very late in their programs, frequently after they have already decided on majors and generated momentum toward specific career plans. One reason for this late introduction is the prerequisite material demanded by our systems courses, typically involving a chain of math, physics, and engineering courses. Using portfolio optimization as the vehicle for introducing systems theory, however, can provide an early introduction to some of the central issues in the field. In particular, the open-loop nature of portfolio optimization simplifies the decision-making context sufficiently to make these problems accessible to younger students. Moreover, the familiarity of financial decision making, regardless of technical background, allows a broad range of students to appreciate the importance and nature of these problems. Here we illustrate these ideas, using portfolio optimization to show how the presence of uncertainty and complexity in decision problems interconnect control, system identification, model reduction, and verification in the design of practical decision systems.**

I. Introduction

Learning platforms are frequently used in education in order to give students a place to experiment with concepts that can be more difficult to learn in a traditional discussion setting. Being able to “try it and see” allows students the opportunity to solve problems and then verify the quality of their solutions. Platforms also play a role in motivating the student, by making learning fun.

For these reasons learning platforms are commonly used in control education. Some common learning platforms include the following: inverted pendula [1], ball and beam systems [2], robotic arms [3], and other mechanical devices. A growing interest in multi-agent systems has likewise motivated team systems such as robot soccer [4] and other “bot” systems [5] that can execute various cooperation strategies to orchestrate efforts to accomplish a common goal. These systems can be powerful platforms for students to solidify their understanding of control and decision processes.

One drawback to these common control platforms is that they require the students to have a mastery of concepts from physics before they can explore the important problems in control. Because of this, students typically are not introduced to systems theory or control before their junior or senior year.

This work is supported by grants from the National Science Foundation (DMS0639328) and the Rollins Center for eBusiness. B. Durtschi and S. Warnick are with the Information and Decision Algorithms Laboratories at Brigham Young University, and M. Skinner is with the Computational Finance program at Carnegie Mellon. Contact durtschi@gmail.com.

Our view is that the central ideas from systems theory should be introduced to students much earlier in their education. This view is motivated by the observation that systems theory and control inform any decision making process, thereby playing a foundational role in a broad range of applications and fields. With earlier exposure, students will be able to decide sooner if they like the study of controls and get a jump start on preparing for the rigors of the field.

To accomplish this, we suggest introducing various learning platforms in areas accessible to younger students. The goal is to simplify the context for discussing central issues in systems theory, providing jumping off points for students to appreciate and explore the depths of the field. Our point of view is that control is fundamentally about making provably good decisions in the face of uncertainty and complexity constraints, and the central problems are 1) the control problem, 2) the system identification problem, 3) the model reduction problem, and 4) the verification problem.

In this paper we introduce portfolio management as a learning platform and illustrate how it can be used to introduce students to each of these four central problems. The area is accessible enough that even students without strong technical backgrounds can appreciate the role of rigor and the importance of systems theory in decision making. Moreover, the area is deep enough that those with strong backgrounds can glean a preview of core research questions involving the interactions between control, system identification, model reduction, and verification problems.

The next section introduces portfolio management as a simple decision problem. Sections three through seven discuss how the portfolio optimization problem guides students through the four previously mentioned problems in control.

II. Portfolio Management

The portfolio management question deals with choosing how to allocate money into different securities with the objective to maximize total wealth at some future time. We have chosen portfolio management as a learning platform because it is conceptually simple and because the objective is clearly parameterized in terms of equity returns. In this way, all questions of information and uncertainty can be posed in terms of what is known about the equity returns. This characterization allows us to reconsider the decision problem repeatedly as we peel back different levels of information and study the impact of uncertainty on our problem.

Moreover, since this problem is open loop, in that investment decisions do not affect future equity returns of the assets, key concepts from systems theory can be introduced without the complexity of feedback interactions. The hope is that students will be motivated to engage the rigors of the discipline necessary to master feedback control if they first appreciate some of the central problems arising from the interaction of uncertainty and complexity in decision problems. Next we describe the portfolio management problem.

Suppose an investor has a choice between holding his money in a risk free cash account with a fixed, positive rate of return, or purchasing any of $n-1$ securities having varying (positive and negative) rates of return. Any of these securities may be purchased at any time, and all that is known about them is their historical price over a finite period of time, $p_i(t),\ i=1,\dots,n$. The goal is to make as much money as possible at time T, by purchasing shares in these securities with a fixed initial investment. The following definitions will help make this objective precise.

A portfolio is a distribution of wealth invested in these assets, characterized by $(s_1(t),\dots,s_n(t))$, where $s_i(t)$ is the number of shares of security i owned at time t. We denote the value of the shares of security i owned at time t by $x_i(t)=p_i(t)s_i(t)$. The value of a portfolio at any particular time is the sum of the value of the securities, $x_1(t)+\dots+x_n(t)$. We let $x_1(t)$ correspond to the value of the risk free cash account. The total return of a security is the price change ratio of the security, given by $r_i(t)=\frac{p_i(t)}{p_i(t-1)}$. This quantity characterizes how the value of a fixed number of shares changes over time. The resulting dynamics of the value of a portfolio over time are given by

$$\begin{bmatrix} x_1(t+1) \\ \vdots \\ x_n(t+1) \end{bmatrix} = \begin{bmatrix} r_1(t+1) & \cdots & 0 \\ 0 & \ddots & 0 \\ 0 & \dots & r_n(t+1) \end{bmatrix} \begin{bmatrix} x_1(t) \\ \vdots \\ x_n(t) \end{bmatrix}. \quad (1)$$

Nevertheless, the investor does not have to keep a fixed number of shares in each security. Instead, s/he can change the distribution of wealth between the securities at each time step. This decision is represented by a set of numbers, $u_i,\ i=1,\dots,n-1$ that indicates the dollar amount that the investor wishes to be moved from the cash account to the ith risky asset. A negative value of u_i represents a dollar amount to be moved from the ith risky asset to the cash account. We will assume there are no transaction costs. The portfolio dynamics incorporating this investor decision then become:

$$x(t+1)=R(t+1)x(t)+R(t+1)Bu(t) \quad (2)$$

where $x(t)\geq 0\ \forall t$, $R(t+1)=diag(r_1(t+1),\dots,r_n(t+1))$,

$$B = \begin{bmatrix} -1 & -1 & \cdots & -1 \\ 1 & 0 & \cdots & 0 \\ 0 & 1 & \ddots & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & \cdots & 1 \end{bmatrix}.$$

The dynamics in (2) describe how the value of a portfolio changes as a function of the investor's reallocation decisions, $u(k)$. In this expression, the return matrix, $R(t+1)$, represents the *future* impact of an external variable over which the investor has no control, while the input, $u(t)$, represents the change in the portfolio distribution over which the investor has complete control. Note that the positivity constraints on x restrict admissible decisions u, allowing the purchase of securities only if you have the money to pay for them. This simple context then motivates a very basic decision problem.

III. The Control Problem

The control problem is to decide which input best improves the performance of the system being controlled. In order to discuss making good decisions we must know the consequences of our decisions and we must have an objective function which ranks the consequences by what is most desirable. Optimal decision making is just computing the choice with the "best" consequence as defined by the objective function. With perfect information about the consequences of our decisions, the control problem becomes a search over possible choices to select the one that best supports the objective. For an in-depth treatment on control see [6].

The control problem naturally surfaces in any discussion of portfolio optimization when we consider the decision to be made by the investor. At some initial time, $t=0$, the investor has an initial amount of cash on hand and no money invested in other securities. Thus his initial portfolio is $x(0)=\begin{bmatrix} x_1(0) & 0 & \dots & 0 \end{bmatrix}^T$. The investor's objective is to allocate his money into different securities at each time step in order to maximize at some future time, T, the total value of the portfolio, $\|x(T)\|_1=x_1(T)+\dots+x_n(T)$. Stated formally,

$$\begin{array}{rrcl} \max\limits_{u_i(1,\dots,T)} & \|x(T)\|_1 & & \\ \text{subject to} & x(t+1) & = & R(t+1)x(t)+R(t+1)Bu(t) \\ & x(0) & = & \begin{bmatrix} x_1(0) & 0 & \dots & 0 \end{bmatrix}' \\ & x(t) & \geq & 0 \quad \forall t \end{array} \quad (3)$$

A student may easily discover that iteratively solving this problem for one time step will yield an optimal solution to the problem for multiple steps. This allows him/her to reduce the problem to a sequence of one step problems.

Example 1: (Perfect Knowledge of Consequences) Consider the problem where $r_i(t),\ i=1,\dots,n$ and $t\in[0,T]$ is given. This problem corresponds to the situation where an investor has perfect knowledge of the returns. To maximize the value of the portfolio one needs to move all the money to the security with the highest return at each time step.

Suppose we can invest in two different securities or keep our money in a cash account. We will start out with \$100, and let x_1 be the amount of money kept in the cash account and x_2 and x_3 be the amount of money invested in the risky securities. Figure 1a shows the value of the two securities over a 110 day period. Figure 1b shows the composition of the optimal portfolio over time as it switches all of the money between the three funds.

Having perfect knowledge about the future has simplified our problem and lets students consider important questions about decision making. For example, can you characterize the nature of optimal solutions to make this computation tractable? How does computational complexity change when

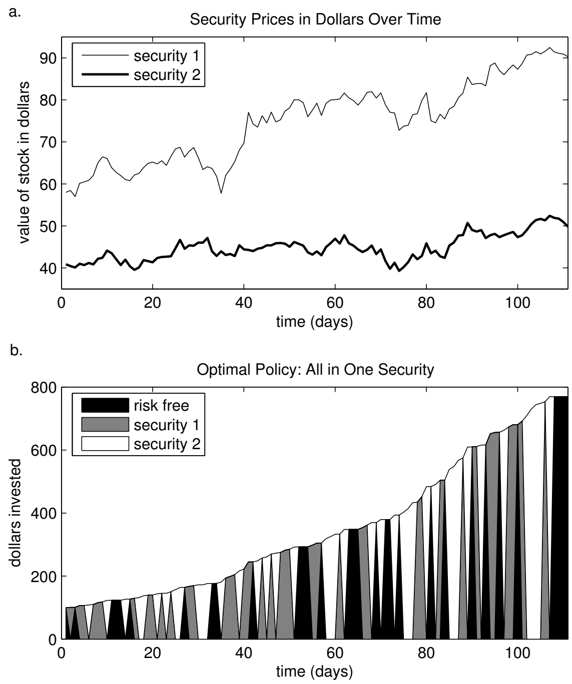

Fig. 1. Prices of two different securities shown above. Assuming perfect knowledge of the future, the optimal policy switches all the money to the security with the highest return at any given time.

considering a sequential decision process where making choices that appear suboptimal now may result in a higher future payoff?

Until now we have assumed perfect knowledge of the future. When we do not have perfect information about the consequences of our actions, we need to estimate a best guess of what the consequences may be in order to make decisions consistent with our objective. A model represents everything we understand about the mapping between choices and consequences. To determine a model for future consequences, students must solve the system identification problem.

IV. The System Identification Problem

Formulating a system identification problem implicitly assumes that something about the mapping of decisions to consequences remains constant over time. This constancy is captured in the choice of model class. The knowledge we have about the true system's behavior is acquired by running experiments that collect input-output data. We use a quality metric to evaluate each model in the class based this data.

Given a class of models, input-output data, and a quality metric, the system identification problem is to select the model from the class which best describes the observed input-output data according to the quality metric, see [7]

A. Sources of Uncertainty

Once a model is selected it becomes the basis for predicting consequences of various decisions. Inaccuracies in the predictions of this chosen model can come from insufficient data, the model class, or the quality metric. These sources of uncertainty are easily seen in portfolio management.

TABLE I

	Actual	Estimated				
sample size		5	15	50	100	1000
mean	1.1	1.155	.989	1.134	1.122	1.088
variance	.1	.040	.148	.143	.109	.102

Example 2: (Uncertainty from Insufficient Data) Suppose we have a risky security. We no longer assume that we have perfect knowledge about the future returns but that they come from a Gaussian distribution. We select our model class to be the class of Gaussian distributions parameterized by mean and variance. Our system identification algorithm is to choose the model whose mean and variance most closely match the sample mean and variance of the historical data.

Suppose the past returns from a risky security are truly generated from a Gaussian distribution with mean, 1.10, and variance, 0.10. By taking historical data as our sample we can compute an estimated mean and variance. As shown in Table I, the estimated mean and variance change depending on how large a sample we use.

This example shows that with finite input-output data the learned model will be different from the true system, which introduces uncertainty into our predictions. Also, as the available historical data increases, the sample mean and variance, and thus the selected model, converge to the same as the true system. This consistency is an indicator of a good system identification algorithm.

Example 3: (Uncertainty Intrinsic to the Model Precision) Considering the previous example, suppose we had enough data that our system identification algorithm could select the correct model from the model class. Would our predictions of future returns necessarily be accurate? In this case, our predictions of future returns are the mean of the Gaussian model we choose as most descriptive of the historical data. However, our model class emphasizes a level of uncertainty in these predictions, characterized by the variance of the Gaussian distribution; we expect our predictions to be accurate, on average, but within a range specified by the variance of our model. Thus, our particular choice of model class builds in an estimate of the level of precision of our predictions, and this precision defines uncertainty intrinsic to our particular model.

Example 4: (Uncertainty Due to Inaccurate Choice of Model Class) Another source of uncertainty is in the selection of the model class itself. While we may choose to represent returns as a Gaussian random variable, it may, in fact, be generated by a completely different process for which a Gaussian process is only an approximation.

Example 5: (Uncertainty Due to the Quality Metric) Suppose we have enough data such that our sample mean is 1.1, but we change our quality metric to choose the model which weights more heavily the most recent week's data. The model selected will not be the same as the actual distribution mean, as our choice of metric reveals our assumptions about the predictive value of the model. Likewise, one may compute a "best fit" between model predictions and data using different norms, and the minimizer in each case may be an entirely different model from the chosen model class.

Frequently the choice of metric is used to guard against uncertainty in the choice of model class. This is accomplished, for example, by ranking each model in the class by the likelihood that it generated the observed data, and then choosing the model that maximizes this likelihood.

Consideration of these uncertainties in modeling leads students to a variety of important questions. How should one characterize uncertainty in a model, and how should this characterization change the control problem formulation? How does one design tests that check whether the underlying assumptions justifying a particular choice of model class, metric, etc. are still consistent with the observed data?

B. Uncertainty's Affect on System ID and Control Problems

Until now we have discussed the control problem and the system ID problem separately. Because of uncertainty, it is interesting to consider how these problems affect one another. Given a predictive model, when does it make sense to treat its predictions as the true future, leaving the formulation of the control problem unchanged? When does it make sense to modify the formulation of the control problem to account for the fact that our estimates of the consequences of various decisions are based on a model, not perfect knowledge? Likewise, should knowing that our choice of a model will be used as the basis for decision making alter the nature of the identification problem? If so, how?

One approach to resolving these issues suggests that both the system identification and control problems should be modified to account for their impact on each other. For example, some control-oriented system identification modifies system identification techniques to be compatible with robust control methodologies [8]. Likewise, robust control can be viewed naturally as identification-oriented control because it modifies classical control techniques to account for explicit uncertainty in the learned model. In the portfolio management example, these issues arise naturally as the objective function in the control problem is modified to account for uncertainty in the predictions based on the identified model. The degree to which the objective function is modified to account for this uncertainty can be scaled by a risk aversion parameter, facilitating an entire class of control problems depending on the degree one is willing to believe the predictions of the identified model.

Example 6: (Control Problem Accounting for Model Uncertainty) Now we assume that we have uncertain predictions of future returns and show how the decision problem is different from the decision problem with perfect knowledge of the future. We let J be our one step decision problem without uncertainty and then consider how uncertainty in the returns affects the solution.

$$\begin{aligned} J = \max_{u_i(1,\dots,T)} \quad & \|x(t+1)\|_1 \\ \text{subject to} \quad x(t+1) &= R(t+1)x(t) + R(t+1)Bu(t) \\ x(t) &\geq 0 \quad \forall t \end{aligned} \tag{4}$$

This optimization requires predictions for the future returns, $r_i(t+1), i = 1,\dots,n$. Assuming that we solve an identification problem to find a Gaussian distribution that best explains the observed data, then we could use the mean for each security as the predicted return for that security. However, this would be assuming that our identified model is perfectly accurate.

One way to formulate the control problem that accounts for our uncertainty in the predictions of our model is to say that each predicted return $\hat{r}_i$ is never more than ε away from the actual return, $|\hat{r}_i - r_i| \leq \varepsilon_i$. This gives us a range of possible values for $J \in [J_{lo}, J_{hi}]$. Now we need to rethink our objective function based on the ranges of J for each possible combination of u_i's. We can think of protecting against the worst case by maximizing the minimum value of J. We could also maximize the maximum, the median, or some other value of J; we could minimize the difference $J_{hi} - J_{lo}$, etc. In each of these situations, considering the sense in which the resulting investment strategy is "best" encourages students to think deeply about the interaction between identification and control.

Another way to formulate the control problem to account for uncertainty in the identified model is to modify the objective to not only maximize predicted returns, but also to minimize the uncertainty intrinsic to these predictions. This is easily accomplished in the case where we model returns as Gaussian random processes, as we can simply modify the original objective to account for the covariance of the random process, given by

$$\max_{u_1(t),\dots,u_{n-1}(t)} \mu_J - \lambda \sigma_J^2. \tag{5}$$

This problem is the celebrated Markowitz model, which is commonly used in portfolio optimization [9] [10].

Yet another approach may discount for uncertainty differently, by only considering the down-side risk of a particular investment. One way of doing this would be to let μ_J be the expected value of J based on the expected returns $\hat{r}_i$, and a fixed portfolio specified by $u(t)$. Then for each previous time, $k = 1,\dots,K$, we compute $y(k) = \mu_J - J(k)$, where $J(k)$ is how well our same portfolio would have performed at time k. $y(k)$ gives us a measure for how much our portfolio would have underperformed our expectation in the past. We then let our uncertainty be the average of the $y(k)$'s for all k. This gives us,

$$\max_{u_1,\dots,u_{n-1}} \mu_J - \lambda \frac{1}{K} \sum_{k=1}^{K} y(k), \tag{6}$$

which essentially yields the Mean Absolute Deviation (MAD) portfolio optimization model [11].

In both the Markowitz and MAD formulations, a risk aversion parameter, λ, describes how much the control objective should accommodate uncertainty in an identified model. By increasing λ we penalize decisions which choose stocks with higher uncertainty.

Using the same three securities that we had in Example 1, but now with uncertainty, Figure 2 shows an optimal policy

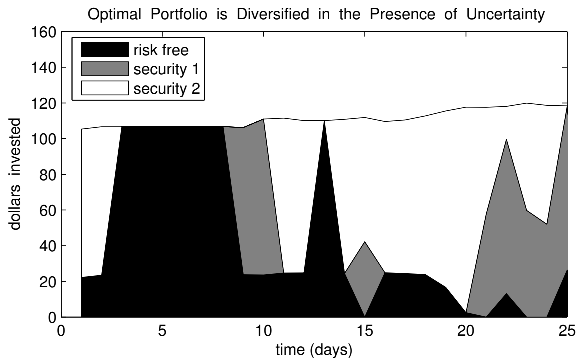

Fig. 2. Using a system identification algorithm to predict future values with uncertainty, yields an optimal risk/reward portfolio that is diversified. These 25 days correspond to the last 25 days from Figure 1.

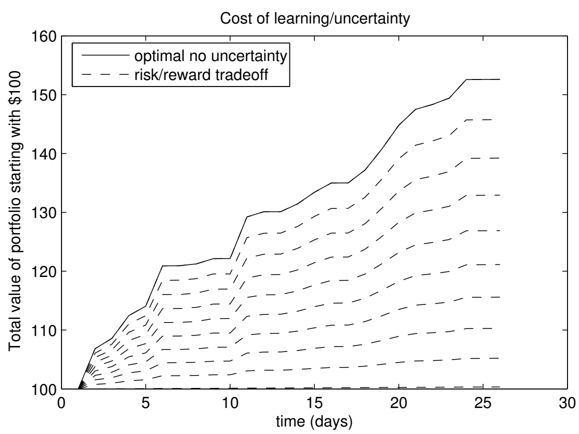

Fig. 3. The cost of uncertainty compared to the optimal strategy. The different dotted lines represent portfolios with different risk tolerances. Because of uncertainty in the future returns, portfolios with lower risk tolerance had to diversify more. Given that the predictions were correct, diversification lowers the optimal performance i.e. robustness incurs a performance cost.

that allows a limited amount of risk. Unlike in Figure 1 when we had perfect knowledge of the future, this portfolio is *diversified*, splitting money between multiple securities to reduce risk. The return from Figure 2 is about 20% compared to about 50% in the optimal portfolio over the same time.

It is also interesting to consider how uncertainty affects the possible solution. For example, suppose that we have an identification algorithm with uncertain predictions that happens to predict the returns exactly. Since we do not know ahead of time that they will be exact predictions, the risk/reward decision strategy will diversify the portfolio more or less depending on the uncertainy in the predictions. Figure 3 shows the loss in returns due to this diversification. In other words, by limiting the incurred risk of a portfolio, the maximum potential is lowered as well. Viewing diversification as a way to overcome uncertainty make portfolio optimization an interesting platform to study control.

V. The Model Reduction Problem

Until now we have assumed that the true system was contained in the model class chosen for system identification. There are however, many problems associated with this assumption. First, it is almost impossible to validate. At most one can only show that the selected model is not invalidated by the data from the true system. Second, models which can actually describe the true system are usually too complex to be efficiently identified or used in computing decisions. Finally, complex models may require more data in order to be identified than is available.

The model reduction problem centers around reducing complexity while retaining as much accuracy as possible. The problem states that given a complex system, G, find a simpler system from a class of simple systems, $\hat{\mathbf{G}}$, such that the difference between them is minimized in some norm.

$$\inf_{\hat{G}\in\hat{\mathbf{G}}} \left\| G - \hat{G} \right\|_n \tag{7}$$

See [6] for more discussion on Model Reduction.

Another way to reduce complexity is to lower computational complexity by decreasing the size of the control and system identification problems. An example of this for portfolio management is to limit the number of stocks being considered. This will avoid introducing the uncertainty that comes from simplifying the model class, however, it may lead to decreased performance because some of the previous options are no longer available for the decision algorithm. The next example shows many ways how approximation may be done by students with portfolio optimization.

Example 7: (Reducing Complexity in Portfolio Choice) Continuing previous examples where we have a simple model for future returns, the Markowitz formulation for the control problem requires a quadratic program to compute the covariance of a portfolio. The solution can not be solved quickly for a large number of securities. The literature is rich with methods for approximating the Markowitz portfolio [12], [13]. One might also consider simplifying the problem by using an alternative formulation such as the MAD model which can be computed with a linear, instead of a quadratic program.

Another alternative for reducing complexity is to limit the number of securities available for selection. One might take the top n performers over the last period of time, or the lowest valued in the Dow Jones Industrial Average, etc. In our example, we limit the number of stocks under consideration from a progressively smaller subset of securities in the American Stock Exchange. Figure 4 shows how the value of a MAD portfolio is affected when choosing subsets of securities of different sizes. In particular showing it is non trivial to restrict stocks to a subset. Choosing a good restrictive subset this figure shows that choosing a subset of securities from which the optimal portfolio performs as well as the non restricted set is non trivial.

The portfolio optimization problem introduces students to consider some important questions in approximation. For example, what approximations can be made that keep solutions close to non-approximated optimal solutions? How does limiting the number of securities affect the computational complexity as well as the performance of the decision algorithm?

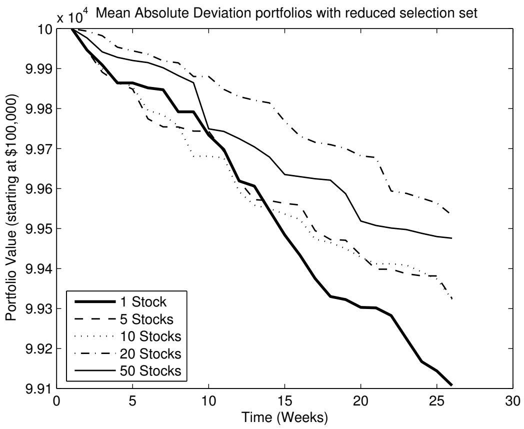

Fig. 4. An experiment of the performance of the MAD model on progressively limited subsets of stocks. Students must answer the question how does limiting the number of stocks affect the performance of the decision algorithm? The downward trend reflects the bear market that has been prevalent over the past year.

VI. The Verification Problem

Once the solution to a control problem has been formulated and has been approximated where necessary to ensure that a solution can actually be computed, some very important questions remain. How do we know that the approximated model is still precise? How do we know that the control solution works as desired? How do we convince non technical people that our methods are sound? These questions motivate the final major problem in control that our platform introduces to students.

The verification stage monitors the performance of a complete solution, including the control, identification, and reduction steps. We call such a complete solution an *algorithmic decision process*. Verification determines whether we observe any new evidence that the assumption justifying many of the choices made in the design of the algorithmic decision process have been violated, thereby motivating a redesign of the solution. Note that we can never prove that a particular design will always work, we simply look for evidence that it begins to fail [14].

Over time verification has become an increasingly important field of research. As controllers are implemented in software and hardware it becomes imperative to verify that these controllers will work as designed. In designing system-on-chip solutions, for example, 70% of the effort is spent on verification [15]. Because of this effort, verification should be considered early-on in the design of solutions to the control, system identification, and model reduction problems.

Example 8: (Verifying Portfolio Optimization) In the portfolio optimization problem we want to use verification to determine whether our decision process is able to select portfolios that are better than the competition. One first attempt can be to run on past data to determine whether the algorithm performes above a specified benchmark.

In addition to verification with past data, running algorithms against others provides a useful means of verification. Many virtual fund managements systems have a way to compare algorithms against each other in hopes of determining which is better. Brigham Young University's Tour de Finance is a platform that allows for student defined competition dynamics as a particular verification method [16].

As students go through the control design process, the Tour de Finance platform gives them a fun competition which will also serve as a verification mechanism. Students can create leagues where algorithms can compete against each other. By seeing which algorithm performs better over time, they can determine which algorithm is better.

VII. Conclusions

We have introduced portfolio management as a learning platform designed to teach students about the decision making process and introduce them to important controls problems earlier in their education. We walked through the example of portfolio management, showing how questions encountered by trying to select an optimal portfolio led to deep questions in the problems of control, system identification, model reduction, and verification. Moreover, classic results in finance such as the Markowitz model lead to exploring the interaction among these problems.

References

[1] M. Yamakita, T. Hoshino, and K. Furuta. Control practice using pendulum. In *American Control Conference*, 1999.

[2] Saul Jimenez and Wen Yu. Stable synchronization control for two ball and bearn systerns. In *Electrical and Electronics Engineering, International Conference on, ICEEE*, 2007.

[3] A. Bindal, S. Mann, B. N. Ahmed, and L. A. Raimundo. An undergraduate system-on-chip (soc) course for computer engineering students. *IEEE Transactions on Education*, May 2005.

[4] R. Hill and A. van den Hengel. Experiences with simulated robot soccer as a teaching tool. In *Information Technology and Applications, International Converence on, ICITA*, 2005.

[5] Salman Ahmed and Mohd N. Karsiti. A testbed for control schemes using multi agent nonholonomic robots. In *Electro/Information Technology, International Conference on*, 2007.

[6] Geir E. Dullerud and Fernando Paganini. *A Course in Robust Control Theory: A Convex Approach*. Springer, 2000.

[7] Lennart Ljung. *System Identification Theory for the User*. PTR Prentice-Hall Inc., 1987.

[8] Jie Chen and Guoxiang Gu. *Control-oriented system identification: An H-inf approach*. Wiley, 2000.

[9] David G. Luenberger. *Investment Science*. Oxford University Press, 1998.

[10] Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, March 1952.

[11] Konno, Hiroshi and Yamazaki, Hiroaki. *Management Science*, May 1991.

[12] W.F. Sharpe. A linear programming approximation for the general portfolio analysis problem. *Journal of Financial and Quantitative Analysis*, December 1971.

[13] B.K. Stone. A linear programming formulation of the general portfolio selection problem. *The Journal of Financial and Quantitative Analysis*, September 1973.

[14] Kameshwar Poolla, Pramod Khargonekar, Ashok Tikku, James Krause, and Krishan Nagpal. A time-domain approach to model validation. *IEEE Transactions on Automatic Control*, May 1994.

[15] Chen Wenwei, Zhang Jinyi, Li Jiao, Ren Xiaojun, and Liu Jiwei. Study on a mixed verification strategy for IP-based soc design. In *High Density Microsystem Design and Packaging and Component Failure Analysis, Conference on*, 2005.

[16] Casey Dougal, David Merriman, Jeffery Humpherys, and Sean Warnick. Competition dynamics in a virtual fund management system: The tour de finance. In *American Control Conference*, 2007.