

Performance Bounds and Suboptimal Policies for Multi-Period Investment

Stephen Boyd

Stanford University
boyd@stanford.edu

Mark T. Mueller

Cambridge, MA
mark.t.mueller@mac.com

Brendan O'Donoghue

Stanford University
bodonoghue85@gmail.com

Yang Wang

Stanford University
yang1024@gmail.com

now

the essence of knowledge

Boston — Delft

Foundations and Trends[®] in Optimization

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

S. Boyd, M. Mueller, B. O'Donoghue, Y. Wang. *Performance Bounds and Suboptimal Policies for Multi-Period Investment*. Foundations and Trends[®] in Optimization, vol. 1, no. 1, pp. 1–72, 2014.

This Foundations and Trends[®] issue was typeset in L^AT_EX using a class file designed by Neal Parikh. Printed on acid-free paper.

ISBN: 978-1-60198-673-3

© 2014 S. Boyd, M. Mueller, B. O'Donoghue, Y. Wang

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

Foundations and Trends[®] in Optimization
Volume 1, Issue 1, 2014
Editorial Board

Editors-in-Chief

Stephen Boyd
Stanford University
United States

Yinyu Ye
Stanford University
United States

Editors

Dimitris Bertsimas
Massachusetts Institute of Technology

John R. Birge
University of Chicago

Robert E. Bixby
Rice University

Emmanuel Candès
Stanford University

David Donoho
Stanford University

Laurent El Ghaoui
University of California, Berkeley

Donald Goldfarb
Columbia University

Michael I. Jordan
University of California, Berkeley

Zhi-Quan (Tom) Luo
University of Minnesota, Twin Cities

George L. Nemhauser
Georgia Institute of Technology

Arkadi Nemirovski
Georgia Institute of Technology

Yurii Nesterov
UC Louvain

Jorge Nocedal
Northwestern University

Pablo A. Parrilo
Massachusetts Institute of Technology

Roman Polyak
George Mason University

Tamás Terlaky
Lehigh University

Michael J. Todd
Cornell University

Kim-Chuan Toh
National University of Singapore

John N. Tsitsiklis
Massachusetts Institute of Technology

Lieven Vandenberghe
University of California, Los Angeles

Robert J. Vanderbei
Princeton University

Stephen J. Wright
University of Wisconsin

Editorial Scope

Topics

Foundations and Trends[®] in Optimization publishes survey and tutorial articles on methods for and applications of mathematical optimization, including the following topics:

- Algorithm design, analysis, and implementation (especially on modern computing platforms)
- Models and modeling systems
- New optimization formulations for practical problems
- Applications of optimization in:
 - Machine learning
 - Statistics
 - Data analysis
 - Signal and image processing
 - Computational economics and finance
 - Engineering design
 - Scheduling and resource allocation
 - and other areas

Information for Librarians

Foundations and Trends[®] in Optimization, 2014, Volume 1, 4 issues. ISSN paper version 2167-3888. ISSN online version 2167-3918. Also available as a combined paper and online subscription.

Foundations and Trends[®] in Optimization
Vol. 1, No. 1 (2014) 1–72
© 2014 S. Boyd, M. Mueller, B. O’Donoghue,
Y. Wang
DOI: 10.1561/2400000001



Performance Bounds and Suboptimal Policies for Multi-Period Investment

Stephen Boyd
Stanford University
boyd@stanford.edu

Mark T. Mueller
Cambridge, MA
mark.t.mueller@mac.com

Brendan O’Donoghue
Stanford University
bodonoghue85@gmail.com

Yang Wang
Stanford University
yang1024@gmail.com

Contents

1	Introduction	2
1.1	Overview	2
1.2	Prior and related work	5
1.3	Outline	8
2	Stochastic Control Formulation	9
2.1	Model	9
2.2	Stage cost function	13
2.3	Post-trade constraints	13
2.4	Transaction and position costs	18
2.5	Quadratic and QP-representable stage cost	21
3	Optimal Policy	22
3.1	Dynamic programming	22
3.2	Quadratic case	23
3.3	No transaction cost case	25
4	Performance Bounds	27
4.1	Bellman inequalities	27
4.2	LMI conditions	28
4.3	Summary	29

5	Approximate Dynamic Programming	30
5.1	Basic idea	30
5.2	Quadratic approximate dynamic programming	31
5.3	ADP based on quadratic underestimators	32
6	Model Predictive Control	33
6.1	Policy	33
6.2	Implementation	34
6.3	Interpretation as an ADP policy	35
6.4	Truncated MPC	36
7	Numerical Examples	38
7.1	Problem data	38
7.2	Computation	40
7.3	Performance bounds and policy performance	42
7.4	Simulation results and trajectories	44
7.5	Robustness to model parameters	46
7.6	Robustness to return distribution	47
8	Conclusions	50
	Appendices	54
A	Expectation of Quadratic Function	55
B	Partial Minimization of Quadratic Function	57
C	\mathcal{S}-Procedure	59
D	LMI Sufficient Condition for Bellman Inequality	61
E	No-Trade Region	63
	References	65

Abstract

We consider dynamic trading of a portfolio of assets in discrete periods over a finite time horizon, with arbitrary time-varying distribution of asset returns. The goal is to maximize the total expected revenue from the portfolio, while respecting constraints on the portfolio such as a required terminal portfolio and leverage and risk limits. The revenue takes into account the gross cash generated in trades, transaction costs, and costs associated with the positions, such as fees for holding short positions. Our model has the form of a stochastic control problem with linear dynamics and convex cost function and constraints. While this problem can be tractably solved in several special cases, such as when all costs are convex quadratic, or when there are no transaction costs, our focus is on the more general case, with nonquadratic cost terms and transaction costs.

We show how to use linear matrix inequality techniques and semidefinite programming to produce a quadratic bound on the value function, which in turn gives a bound on the optimal performance. This performance bound can be used to judge the performance obtained by any suboptimal policy. As a by-product of the performance bound computation, we obtain an approximate dynamic programming policy that requires the solution of a convex optimization problem, often a quadratic program, to determine the trades to carry out in each step. While we have no theoretical guarantee that the performance of our suboptimal policy is always near the performance bound (which would imply that it is nearly optimal) we observe that in numerical examples the two values are typically close.

1

Introduction

1.1 Overview

In this paper we formulate the discrete-time finite horizon time-varying multi-period investment problem as a stochastic control problem. By using state variables that track the *value* of the assets, instead of more traditional choices of states such as the number of shares or the fraction of total value, the stochastic control problem has *linear* (but random) dynamics. Assuming that the costs and constraints are convex, we arrive at a linear convex stochastic control problem.

This problem can be effectively solved in two broad cases. When there are no transaction costs, the multi-period investment problem can be reduced to solving a set of standard single-period investment problems; the optimal policy in this case is to simply rebalance the portfolio to a pre-computed optimal portfolio in each step. Another case in which the problem can be effectively solved is when the costs are quadratic and the only constraints are linear equality constraints. In this case standard dynamic programming (DP) techniques can be used to compute the optimal trading policies, which are affine functions of the current portfolio. We describe these special cases in more detail in §3.3 and §3.2. The problem is also tractable when the number of assets

is very small, say two or three, in which case brute force (numerical) dynamic programming can be used to compute an optimal policy.

Most problems of interest, however, include significant transaction costs, or include terms that are not well approximated by quadratic functions. In these cases, the optimal investment policy cannot be tractably computed. In such situations, several approaches can be used to find suboptimal policies, including approximate dynamic programming (ADP) and model predictive control (MPC). The performance of any suboptimal policy can be evaluated using Monte Carlo analysis, by simulation over many return trajectories. An obvious practical (and theoretical) question is, how suboptimal is the policy? In this paper we address this question.

Using linear matrix inequality (LMI) techniques widely used in control system analysis and design [18, 35, 80], we construct a (numerical) bound on the best performance that can be attained, for a given problem. The method requires the construction and solution of a semidefinite program (SDP), a convex optimization problem involving matrix inequalities. We can compare the bound on performance with the performance attained by any suboptimal policy; when they are close, we conclude that the policy is approximately optimal (and that the performance bound is nearly tight). Even when the performance bound and suboptimal policy performance are not close, we at least have a bound on how suboptimal our suboptimal policy can be.

The performance bound computation yields a quadratic approximation (in fact, underestimator) of the value functions for the stochastic control problem. These quadratic value function approximations can be used in an ADP policy, or as the terminal cost in an MPC policy. While we have no a priori guarantee that the gap between the performance bound and the performance of the ADP policy will always be small, simulations show that the ADP and MPC policies achieve performance that is often nearly optimal.

Our methods for computing the performance bound, as well as implementing the ADP and MPC suboptimal policies, rely on (numerically) solving convex optimization problems, for which there are efficient and reliable algorithms available [20, 72, 77, 73, 102]. The per-

formance bound computation requires solving SDPs [20, 95], which can be done using modern interior-point cone solvers such as SeDuMi or SDPT3 [88, 92, 94]. Parser-solvers such as CVX or YALMIP [40, 58] allow the user to specify the SDPs in a natural high-level mathematical description form, greatly reducing the time required to form and solve the SDPs. The SDPs that we solve involve T matrices of size $n \times n$, where n is the number of assets, and T is the trading period horizon. These SDPs can be challenging to solve (depending on n and T , of course), using generic methods; but this computation is done once, off-line, before trading begins.

Evaluating the ADP suboptimal policy in each period (*i.e.*, determining the trades to execute) requires solving a small and structured convex optimization problem with (on the order of) n scalar variables. Solving these problems using generic solvers might take seconds, or even minutes, depending on the problem size and types of constraints and objective terms. But recent advances have shown that if the solver is customized for the particular problem family, orders of magnitude speed up is possible [64, 65, 63, 62, 99]. This means that the ADP trading policies we design can be executed at time scales measured in milliseconds or microseconds for modest size problems (say, tens of assets), even with complex constraints. In addition, the trading policies we design can be tested and verified via Monte Carlo simulation very efficiently. For example, the simulation of the numerical examples of the ADP policies reported in this paper required the solution of around 50 million quadratic programs (QPs). These were solved in a few hours on a desktop computer using custom solvers generated by CVXGEN, a code generator for embedded convex optimization [64].

Evaluating the MPC policy also requires the solution of a structured convex optimization problem, with (on the order of) nT variables. If a custom solver is used, the computational effort required is approximately T times the effort required to evaluate the ADP policy. One major advantage of MPC is that it does not require any pre-computation; to implement the ADP policy, we must first solve a large SDP to find the approximate value functions. MPC can thus directly incorporate real-time signals such as changes in future return statistics.

1.2 Prior and related work

Portfolio optimization has been studied and used for more than 60 years. In this section our goal is to give a brief overview of some of the important research in this area, focussing on work related to our approach. Readers interested in a broader overview of the applications of stochastic control and optimization to economics and finance should refer to, *e.g.*, [1, 34, 45, 76, 90, 104].

Single-period portfolio optimization

Portfolio optimization was introduced by Markowitz in 1952 [61]. He formulated a single period portfolio investment problem as a quadratic optimization problem with an objective that trades off expected return and variance. Since this first work, many papers have extended the single period portfolio optimization framework. For example, Goldsmith [38] is one of the first papers to include an analysis of the effect of transaction costs on portfolio selection. Modern convex optimization methods, such as second-order cone programming (SOCP), are applied to portfolio problems with transaction costs in [57, 56]. Convex optimization methods have also been used to handle more sophisticated measures of risk, such as conditional value at risk (CVaR) [82, 50].

Dynamic multi-period portfolio optimization

Early attempts to extend the return-variance trade-off to multi-period portfolio optimization include [91, 71]. One of the first works on multi-period portfolio investment in a dynamic programming framework is by Merton [68]. In this seminal paper, the author considers a problem with one risky asset and one risk-free asset; at each continuous time instant, the investor chooses what proportion of his wealth to invest and what to consume, seeking to maximize the total utility of the wealth consumed over a finite time horizon. When there are no constraints or transaction costs, and under some additional assumptions on the investor utility function, Merton derived a simple closed-form expression for the optimal policy. In a companion paper [84], Samuelson derived the discrete-time analog of Merton's approach.

Constantinides [26] extended Samuelson's discrete-time formulation to problems with proportional transaction costs. In his paper, Constantinides demonstrated the presence of a convex 'no-trade cone'. When the portfolio is within the cone the optimal policy is not to trade; outside the cone, the optimal policy is to trade to the boundary of the cone. (We will see that the policies we derive in this paper have similar properties.) Davis and Norman [29] and Dumas and Luciano [33] derived similar results for the continuous-time formulation. In [28], the authors consider a specific multi-period portfolio problem in continuous time, where they derive a formula for the minimum wealth needed to hedge an arbitrary contingent claim with proportional transaction costs. More recent work includes [93, 23, 24]; in these the authors develop affine recourse policies for discrete time portfolio optimization.

Log-optimal investment

A different formulation for the multi-period problem was developed by Kelly [49], where it was shown that a log-optimal investment strategy maximizes the long-term growth rate of cumulative wealth in horse-race markets. This was extended in [21] to general asset returns and further extended to include all frictionless stationary ergodic markets in [3] and [27]. More recently, Iyengar [44] extended these problems to include proportional transaction costs.

Linear-quadratic multi-period portfolio optimization

Optimal policies for unconstrained linear-quadratic portfolio problems have been derived for continuous-time formulations by Zhou and Li [103], where the authors solve a continuous-time Riccati equation to compute the value function. In [53] this was extended to include a long-only constraint. Skaf and Boyd [87], and Gârleanu and Pederson [37], point out that the multi-period portfolio optimization problem with linear dynamics and convex quadratic objective can be solved exactly. For problems with more complex objective terms, such as proportional transaction costs, Skaf and Boyd use the value functions for an associated quadratic problem as the approximate value functions in an ADP

policy. In [43] the authors formulate a multi-period portfolio problem as a linear stochastic control problem, and propose an MPC policy.

Optimal execution

An important special case of the multi-period problem is the optimal execution problem, where we seek to execute a large block of trades while incurring as small a cost as possible. Bertsimas and Lo [16] model price impact, in which trading affects the asset prices, and derive an optimal trading policy using dynamic programming methods. Almgren and Chriss [4] address the optimal execution problem, including volatility of revenue. They show that the optimal policy can be obtained with additional restrictions on the price dynamics.

Performance bounds

In problems for which an optimal policy can be found, the optimal performance serves as a (tight) bound on performance. The present paper focuses on developing a numerical bound on the optimal performance for problems for which the optimal policy cannot be found.

Brown and Smith [22] compute a bound on optimal performance and derive a heuristic policy that achieves performance close to the bound. Their bound is given by the performance of an investor with perfect information about future returns, plus a clairvoyance penalty.

In [41], the authors construct an upper bound on a continuous time portfolio utility maximization problem with position limits. They do this by solving an unconstrained ‘fictitious problem’ which provides an upper bound on the value function of the original problem.

In [70], the authors describe a class of linear rebalancing policies for the discrete-time portfolio optimization problem. They develop several bounds, including a bound based on a clairvoyant investor and a bound obtained by solving an unconstrained quadratic problem.

Desai et al. [32] develop a bound for an optimal stopping problem, which is useful in a financial context for the pricing of American or Bermudan derivatives amongst other applications. The bound is derived from a dual characterization of optimal stopping problems as optimization problems over the space of martingales.

1.3 Outline

We structure our paper as follows. In chapter 2 we formulate a general multi-period investment problem as a linear convex stochastic control problem, using somewhat nontraditional state variables, and give examples of (convex) stage cost terms and portfolio constraints that arise in practical investment problems, as well as mentioning some nonconvex terms and constraints that do not fit our model. In chapter 3 we review the dynamic programming solution of the stochastic control problem, including the special case when the stage costs are convex quadratic. In chapter 4 we give our method for finding a performance bound in outline form; the full derivations are pushed to appendices A–C. We describe MPC in chapter 6. In chapter 7 we report numerical results for several examples, using both ADP and MPC trading policies.

References

- [1] J. Adda and R. Cooper. *Dynamic economics: Quantitative methods and applications*. MIT Press, 2003.
- [2] M. Åkerblad and A. Hansson. Efficient solution of second order cone program for model predictive control. *International Journal of Control*, 77(1):55–77, January 2004.
- [3] P. Algoet and T. Cover. Asymptotic optimality and asymptotic equipartition properties of log-optimum investment. *The Annals of Probability*, 16(2):876–898, April 1988.
- [4] R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3(2):5–39, December 2001.
- [5] S. Basak and A. Shapiro. Value-at-risk-based risk management: Optimal policies and asset prices. *Review of Financial Studies*, 14(2):371–405, February 2001.
- [6] R. Bellman. *Dynamic Programming*. Dover Publications, 1957.
- [7] A. Bemporad and C. Filippi. Suboptimal explicit receding horizon control via approximate multiparametric quadratic programming. *Journal of Optimization Theory and Applications*, 117(1):9–38, November 2004.
- [8] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos. The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1):3–20, 2002.
- [9] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust optimization*. Princeton University Press, 2009.

- [10] D. Bertsekas. *Dynamic Programming and Optimal Control: Volume 1*. Athena Scientific, 2005.
- [11] D. Bertsekas. Dynamic programming and suboptimal control: A survey from ADP to MPC. *European Journal of Control*, 11(4-5):310–334, October 2005.
- [12] D. Bertsekas. Rollout algorithms for constrained dynamic programming. Technical report, Laboratory for Information and Decision Systems, MIT, 2005.
- [13] D. Bertsekas. *Dynamic Programming and Optimal Control: Volume 2*. Athena Scientific, 2007.
- [14] D. Bertsekas and S. Shreve. *Stochastic optimal control: The discrete-time case*. Athena Scientific, 1996.
- [15] D. Bertsekas and J. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- [16] D. Bertsimas and A. Lo. Optimal control of execution costs. *Journal of Financial Markets*, 1(1):1–50, April 1998.
- [17] J. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer Series in Operations Research and Financial Engineering. Springer, 1997.
- [18] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. Society for Industrial and Applied Mathematics, 1994.
- [19] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- [20] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [21] L. Breiman. Optimal gambling systems for favorable games. In *Proc. 4th Berkeley Symp. Math. Statist. Probab.*, volume 1, pages 65–78, 1961.
- [22] D. Brown and J. Smith. Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. *Management Science*, 57(10):1752–1770, October 2011.
- [23] G. Calafiore. Multi-period portfolio optimization with linear control policies. *Automatica*, 44(10):2463–2473, October 2008.

- [24] G. Calafiore. An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM J. Control Optim.*, 48(4):2254–2274, June 2009.
- [25] G. Calafiore. Random convex programs. *SIAM J. Optim.*, 20(6):3427–3464, December 2010.
- [26] G. Constantinides. Multiperiod consumption and investment behavior with convex transaction costs. *Management Science*, 25(11):1127–1137, November 1979.
- [27] T. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, January 1991.
- [28] J. Cvitanić and I. Karatzas. Hedging and portfolio optimization under transaction costs: A Martingale approach. *Mathematical Finance*, 6(2):133–165, April 1996.
- [29] M. Davis and A. Norman. Portfolio selection with transaction costs. *Mathematics of Operations Research*, 15(4):676–713, November 1990.
- [30] D. De Farias and B. Van Roy. The linear programming approach to approximate dynamic programming. *Operations Research*, 51(6):850–865, November 2003.
- [31] E. Denardo. *Dynamic Programming: Models and Applications*. Prentice-Hall, 1982.
- [32] V. Desai, V. Farias, and C. Moallemi. Pathwise optimization for optimal stopping problems. *Management Science*, June 2012.
- [33] B. Dumas and E. Luciano. An exact solution to a dynamic portfolio choice problem under transaction costs. *The Journal of Finance*, 46(2):577–595, June 1991.
- [34] J. Dupačová, J. Hurt, and J. Štěpán. *Stochastic modeling in economics and finance*. Applied optimization. Kluwer Academic Publishers, 2002.
- [35] L. El Ghaoui and S. Niculescu. *Advances in linear matrix inequality methods in control*. Advances in design and control. SIAM, 2000.
- [36] C. Garcia, D. Prett, and M. Morari. Model predictive control: Theory and practice. *Automatica*, 25(3):335–348, May 1989.
- [37] N. Gârleanu and L. Pedersen. Dynamic trading with predictable returns and transaction costs. Manuscript, available at <http://pages.stern.nyu.edu/~lpederse/papers/DynamicTrading.pdf>, 2011.
- [38] D. Goldsmith. Transactions costs and the theory of portfolio selection. *Journal of Finance*, 31(4):1127–39, September 1976.

- [39] G. Goodwin, M. Seron, and J. De Doná. *Constrained control and estimation*. Springer, 2005.
- [40] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. <http://cvxr.com/cvx>, April 2011.
- [41] M. Haugh, L. Kogan, and J. Wang. Evaluating portfolio policies: A duality approach. *Operations Research*, 54(3):405–418, June 2006.
- [42] O. Hernandez-Lerma and J. Lasserre. Linear programming approximations for Markov control processes in metric spaces. *Acta Applicandae Mathematicae*, 51:123–139, 1998.
- [43] F. Herzog, G. Dondi, and H. Geering. Stochastic model predictive control and portfolio optimization. *International Journal of Theoretical and Applied Finance*, 10(2):203–233, 2007.
- [44] G. Iyengar and T. Cover. Growth optimal investment in horse race markets with costs. *IEEE Transactions on Information Theory*, 46(7):2675–2683, November 2000.
- [45] K. Judd. *Numerical Methods in Economics*. The MIT Press, 1998.
- [46] Y. Kabanov. Hedging and liquidation under transaction costs in currency markets. *Finance and Stochastics*, 3(2):237–248, 1999.
- [47] Y. Kabanov, M. Rásonyi, and C. Stricker. On the closedness of sums of convex cones in L^0 and the robust no-arbitrage property. *Finance and Stochastics*, 7(3):403–411, 2003.
- [48] Y. Kabanov and C. Stricker. The Harrison–Pliska arbitrage pricing theorem under transaction costs. *Journal of Mathematical Economics*, 35(2):185–196, 2001.
- [49] J. Kelly. A new interpretation of information rate. *IRE Transactions on Information Theory*, 2(3):185–189, September 1956.
- [50] P. Krokmal, J. Palmquist, and S. Uryasev. Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of Risk*, 4(2):11–27, 2002.
- [51] M. Kvasnica, P. Grieder, and M. Baotić. Multi-Parametric Toolbox (MPT), 2004.
- [52] W. Kwon and S. Han. *Receding Horizon Control*. Springer-Verlag, 2005.
- [53] X. Li, X. Zhou, and A. Lim. Dynamic mean-variance portfolio selection with no-shorting constraints. *SIAM J. Control Optim.*, 40(5):1540–1555, January 2002.

- [54] B. Lincoln and A. Rantzer. Relaxing dynamic programming. *IEEE Transactions on Automatic Control*, 51(8):1249–1260, August 2006.
- [55] A. Lo and M. Mueller. Warning: Physics envy may be hazardous to your wealth!, 2010.
- [56] M. Lobo, M. Fazel, and S. Boyd. Portfolio optimization with linear and fixed transaction costs. *Annals of Operations Research*, 152(1):341–365, July 2007.
- [57] M. Lobo, L. Vandenberghe, S. Boyd, and H. Lebert. Applications of second-order cone programming. *Linear Algebra and its Applications*, 284:193–228, November 1998.
- [58] J. Löfberg. YALMIP: A toolbox for modeling and optimization in Matlab. In *Proceedings of the CACSD Conference*, 2004.
- [59] J. Maciejowski. *Predictive Control with Constraints*. Prentice-Hall, 2002.
- [60] A. Manne. Linear programming and sequential decisions. *Management Science*, 6(3):259–267, April 1960.
- [61] H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, March 1952.
- [62] J. Mattingley and S. Boyd. Real-time convex optimization in signal processing. *IEEE Signal Processing Magazine*, 23(3):50–61, June 2009.
- [63] J. Mattingley and S. Boyd. Automatic code generation for real-time convex optimization. In D. P. Palomar and Y. C. Eldar, editors, *Convex optimization in signal processing and communications*, pages 1–41. Cambridge University Press, 2010.
- [64] J. Mattingley and S. Boyd. CVXGEN: A code generator for embedded convex optimization. *Optimization and Engineering*, 13(1):1–27, 2012.
- [65] J. Mattingley, Y. Wang, and S. Boyd. Code generation for receding horizon control. In *IEEE Multi-Conference on Systems and Control*, pages 985–992, 2010.
- [66] J. Mattingley, Y. Wang, and S. Boyd. Receding horizon control: Automatic generation of high-speed solvers. *IEEE Control Systems Magazine*, 31(3):52–65, June 2011.
- [67] D. Mayne, J. Rawlings, C. Rao, and P. Sokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, June 2000.

- [68] R. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, 51(3):247–257, August 1969.
- [69] S. Meyn and R. Tweedie. *Markov chains and stochastic stability*. Communications and control engineering. Cambridge University Press, 2009.
- [70] C. Moallemi and M. Sağlam. Dynamic portfolio choice with linear rebalancing rules. Available at <http://www.columbia.edu/~ms3760/linear.pdf>, 2012.
- [71] J. Mossin. Optimal multiperiod portfolio policies. *The Journal of Business*, 41(2):215–229, April 1968.
- [72] Y. Nesterov and A. Nemirovsky. *Interior-Point Polynomial Methods in Convex Programming*. SIAM, 1994.
- [73] J. Nocedal and S. Wright. *Numerical Optimization*. Springer, 1999.
- [74] B. O’Donoghue, G. Stathopoulos, and S. Boyd. A splitting method for optimal control. *IEEE Trans. Autom. Control*, 2013. To appear.
- [75] B. O’Donoghue, Y. Wang, and S. Boyd. Min-max approximate dynamic programming. In *Proceedings IEEE Multi-Conference on Systems and Control*, pages 424–431, September 2011.
- [76] H. Pham. *Continuous-time stochastic control and optimization with financial applications*. Springer Series in Stochastic modelling and applied probability. Springer, 2009.
- [77] F. Potra and S. Wright. Interior-point methods. *Journal of Computational and Applied Mathematics*, 124(1–2):281–302, 2000.
- [78] W. Powell. *Approximate dynamic programming: Solving the curses of dimensionality*. John Wiley & Sons, 2007.
- [79] A. Prékopa. *Stochastic programming*. Mathematics and its applications. Kluwer Academic Publishers, 1995.
- [80] M. Rami and L. El Ghaoui. LMI optimization for nonstandard Riccati equations arising in stochastic control. *IEEE Trans. Autom. Control*, 41(11):1666–1671, November 1996.
- [81] C. V. Rao, S. J. Wright, and J. B. Rawlings. Application of interior point methods to model predictive control. *Journal of optimization theory and applications*, 99(3):723–757, November 2004.
- [82] R. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *The Journal of Risk*, 2(3):21–42, 2000.

- [83] S. Ross. *Introduction to Stochastic Dynamic Programming: Probability and Mathematical*. Academic Press, 1983.
- [84] P. Samuelson. Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics*, 51(3):239–246, August 1969.
- [85] W. Schachermayer. The fundamental theorem of asset pricing under proportional transaction costs in finite discrete time. *Mathematical Finance*, 14(1):19–48, January 2004.
- [86] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on stochastic programming: Modeling and theory*. MPS-SIAM series on optimization. Society for Industrial and Applied Mathematics, 2009.
- [87] J. Skaf and S. Boyd. Multi-period portfolio optimization with constraints and transaction costs. http://www.stanford.edu/~boyd/papers/dyn_port_opt.html, 2008. Manuscript.
- [88] J. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11:625–653, 1999. Software available at <http://sedumi.ie.lehigh.edu/>.
- [89] M. Sznaier, R. Suarez, and J. Cloutier. Suboptimal control of constrained nonlinear systems via receding horizon constrained control Lyapunov functions. *International Journal on Robust and Nonlinear Control*, 13(3–4):247–259, March 2003.
- [90] C. Tapiero. *Applied stochastic models and control for finance and insurance*. Kluwer, 1998.
- [91] J. Tobin. The theory of portfolio selection. In F. Hahn and F. Brechling, editors, *The Theory of Interest Rates*. Macmillan, 1965.
- [92] K. Toh, M. Todd, and R. Tütüncü. SDPT3—A Matlab software package for semidefinite programming, version 1.3. *Optimization Methods and Software*, 11(1):545–581, 1999.
- [93] U. Topcu, G. Calafiore, and L. El Ghaoui. Multistage investments with recourse: A single-asset case with transaction costs. In *Proceedings of the 47th Conference on Decision and Control*, pages 2398–2403, Cancun, Mexico, 2008.
- [94] R. Tütüncü, K. Toh, and M. Todd. Solving semidefinite-quadratic-linear programs using SDPT3. *Mathematical Programming*, 95(2):189–217, 2003.
- [95] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.

- [96] Y. Wang and S. Boyd. Performance bounds for linear stochastic control. *Systems & Control Letters*, 58(3):178–182, March 2009.
- [97] Y. Wang and S. Boyd. Approximate dynamic programming via iterated Bellman inequalities. http://www.stanford.edu/~boyd/papers/adp_iter_bellman.html, 2010. Manuscript.
- [98] Y. Wang and S. Boyd. Fast model predictive control using online optimization. *IEEE Transactions on Control Systems Technology*, 18(2):267–278, March 2010.
- [99] Y. Wang and S. Boyd. Fast evaluation of quadratic control-Lyapunov policy. *IEEE Transactions on Control Systems Technology*, 19(4):939–946, July 2011.
- [100] Y. Wang and S. Boyd. Performance bounds and suboptimal policies for linear stochastic control via LMIs. *International Journal of Robust and Nonlinear Control*, 21(14):1710–1728, September 2011.
- [101] P. Whittle. *Optimization Over Time: Dynamic Programming and Stochastic Control*. John Wiley & Sons, 1982.
- [102] S. Wright. *Primal-Dual Interior-Point Methods*. SIAM, 1997.
- [103] X. Zhou and D. Li. Continuous-time mean-variance portfolio selection: A stochastic LQ framework. *Applied Mathematics & Optimization*, 42(1):19–33, 2000.
- [104] W. Ziemba and R. Vickson. *Stochastic Optimization Models in Finance*. World Scientific, 2006.