



Coherent Risk-Adjusted Decisions over Time: a Bilevel Programming Approach (Renewal)

Jonathan Eckstein
RUTGERS, THE STATE UNIVERSITY OF NEW JERSEY

10/17/2019
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

DISTRIBUTION A: Distribution approved for public release.

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Executive Services, Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.					
1. REPORT DATE (DD-MM-YYYY) 11-11-2019		2. REPORT TYPE Final Performance		3. DATES COVERED (From - To) 15 Jul 2015 to 14 Jul 2019	
4. TITLE AND SUBTITLE Coherent Risk-Adjusted Decisions over Time: a Bilevel Programming Approach (Renewal)				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER FA9550-15-1-0251	
				5c. PROGRAM ELEMENT NUMBER 61102F	
6. AUTHOR(S) Jonathan Eckstein				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) RUTGERS, THE STATE UNIVERSITY OF NEW JERSEY 123 WASHINGTON ST STE 510 NEWARK, NJ 7102-3026 US				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AF Office of Scientific Research 875 N. Randolph St. Room 3112 Arlington, VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR RTA2	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-AFOSR-VA-TR-2019-0340	
12. DISTRIBUTION/AVAILABILITY STATEMENT A DISTRIBUTION UNLIMITED: PB Public Release					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This report details the research performed under AFOSR grant FA9550-15-1-0251. We investigated and published peerreviewed papers concerning the following topics: risk-averse control of various kinds of stochastic systems, risk measurement of partially observable Markov processes, fundamental modeling issues risk-averse optimization over time, numerical methods for large-scale non-differentiable optimization, and application of operations-research-related optimization techniques to problems in machine learning.					
15. SUBJECT TERMS Coherent Risk-Adjusted Decision, Bilevel Programming					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON FAHROO, FARIBA
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (Include area code) 703-696-8429

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

DISTRIBUTION A: Distribution approved for public release.

Coherent Risk-Adjusted Decisions over Time: a Bilevel Programming Approach (Renewal)

Final Performance Report: AFOSR Grant FA9550-15-1-0251

September 2019

Jonathan Eckstein (PI) and Andrzej Ruszczyński (Co-PI)

In this project, we investigated a large number of topics related to risk-averse and large-scale optimization. We summarize our accomplishments below, categorized by the topic area.

Risk-averse control of discrete-time systems

We extended our earlier theory of risk-averse control of Markov systems to discrete-time systems with history-dependent transitions. We developed the concept of a process-based risk measure, which allows for risk modeling in such systems with a collection of static risk measures, indexed by the process history. We develop associated dynamic programming equations for finite-horizon problems. The results are included in a paper:

J. Fan, A. Ruszczyński, “Process-Based Risk Measures and Risk-Averse Control of Discrete-Time Systems,” *Mathematical Programming*, available online November 2018.

This research was part of the PhD dissertation of Jingnan Fan, defended in December 2017.

Risk-averse control of partially observable discrete-time Markov systems

We have continued research on risk models and risk-averse control for Markov systems in which part of the state vector is not directly observable. For such systems, we developed the theory of time consistency, introduced Markov risk measures, and developed dynamic programming equations. The results are described in the papers:

J. Fan, A. Ruszczyński, “Dynamic Risk Measures for Finite-State Partially Observable Markov Decision Problems,” *Proceedings of the Conference on Control and its Applications*, Paris 2015, pp. 153-158.

J. Fan, A. Ruszczyński, “Risk measurement and risk-averse control of partially observable discrete-time Markov systems,” *Mathematical Methods of Operations Research* 88(2):161-184 (2018).

The research was part of the PhD dissertation of Jingnan Fan, defended in December 2017.

This research direction led to the discovery of new risk models, which we call *risk forms*, that allow dependence on a varying probability distribution. Such forms are needed in partially observable systems because of the Bayesian updates of belief states. The theory of risk forms has been initiated in the paper:

D. Dentcheva, A. Ruszczyński, “Risk forms: representation, disintegration, and application to partially observable two-stage systems,” *Mathematical Programming* available online February 2019.

Risk-averse control of continuous-time Markov chains

We introduced Markov risk measures for continuous-time Markov systems with a discrete state space. By using techniques of set-valued analysis, we developed a differential representation of such risk measures, generalizing the classical backward Kolmogorov equations for the expected value case. Next, we analyzed the optimal control problem for such systems and developed optimality conditions for it. Finally, we developed tight discrete-time approximations for the problem. The results have been presented in two publications:

D. Dentcheva, A. Ruszczyński, "Time-Consistent Risk Measures for Continuous-Time Markov Chains," *SIAM Journal on Financial Math* 9(2):690-715 (2018).

D. Dentcheva, A. Ruszczyński, "Risk-Averse Control of Continuous-Time Markov Chains," *Proceedings of SIAM Conference on Control*, Pittsburgh, PA 2017.

Risk measurement in partially observable continuous-time Markov systems

We considered the situation of a continuous-time Markov chain whose state is not observed, and an observation process whose dynamics depends on the state of the chain. We developed a theory of risk filters for this model, and have shown that they can be described by a system of forward-backward stochastic differential equations. This research was part of the PhD dissertation of Ruofan Yan, defended in December 2017.

Modeling risk-averse optimization over time

In problems in which uncertainty is resolved at more than one point in time, interleaved with decisions, coherent risk measures pose various time-consistency conundrums. Time consistency is violated when the later portions of a purportedly optimal sequence of decisions may appear suboptimal from the viewpoint of a later time. The standard formulation approach is to use dynamic (multi-period) risk measures constructed in a nested manner such that time-inconsistent solutions are impossible. However, these nested risk measures have some other properties that can make them unnatural or difficult for decision makers to interpret.

In one published work, we consider techniques closely approximating a time-inconsistent risk measure by a time-consistent one. This technique allows the decision maker to use a time consistent-objective function that is as close as possible to an otherwise more intuitive time-inconsistent one. This work appears as

T. Asamov, A. Ruszczyński, "Time-consistent approximations of risk-averse multistage stochastic optimization problems," *Mathematical Programming* 153 (2), 2015, 459-493.

We also developed a general modeling technique in which multilevel optimization constraints impose time consistency, allowing the formulator to use any desired objective function without violating time consistency. We have made a lengthy theoretical study of this modeling technique, showing under which conditions it is equivalent to prior approaches, and when it is different. We also showed that in various simple situations, it leads to $\mathcal{NP}$ -hard models. In a simple supply chain optimization setting, we studied such models and showed that the resulting optimization models were practically solvable by a modern integer programming solver (Gurobi), producing solutions that were sometimes very different from those obtained using nested time-consistent objective formulation techniques. This work was published in:

J. Eckstein, D. Eskandani, and J. Fan. “Multilevel Optimization Modeling for Risk-Averse Stochastic Programming,” *INFORMS Journal on Computing* 28(1):112-128 (2016; with online supplements containing proofs etc).

We then continued work on this area, exploring the effect of similar formulation techniques on portfolio rebalancing and hydropower planning problems, using various risk measures. In some cases, such as portfolio rebalancing, the multilevel formulations have solutions nearly identical to those arising with other techniques, but in other cases the solutions are quite different. This empirical work is not yet published in a journal, but will be part of the dissertation of Deniz Eskandani, scheduled for completion in December 2019.

More recently, Deniz has been working on the practical solution of stochastic bilevel linear programming problems, in which a “leader” makes an initial decision (subject to linear constraints), and then a random event is realized (typically from a finite distribution). A “follower” then takes a decision subject to constraints that are a function of the leader decision and the random event. The goal is for the leader to maximize the expected value of some (usually linear) function of the leader and follower variables. Such models include as special cases three-stage risk-averse models formulated with our multilevel modeling techniques.

We are presently finalizing work on efficiently computing approximate solutions of such problems. The straightforward method of converting the problem to a disjunctive program and solving it with a standard MILP solver works for smaller problem instances, but breaks down for larger ones, as the size of the search tree becomes unmanageable. After considerable experimentation with various relaxations (including those using semidefinite programming), we have settled on a procedure employing a specialized form of bundle method, in particular the method described in the working paper

W. de Oliveira and J. Eckstein, “A Bundle Method for Exploiting Additive Structure in Difficult Optimization Problems”, E-print 2015-05-4935, Optimization Online, May 2015.

For most problems, Deniz has been able to use such a method, in conjunction with a specially formulated heuristic, to find very close lower and upper bounds on the optimal solution value of stochastic bilevel programs. A few details relating to how to deal with unbounded subproblems remain to be dealt with, after which Deniz should be wrapping up her dissertation.

Numerical methods for large-scale non-differentiable optimization

In a related line of research, we developed new large scale non-differentiable optimization methods applicable to a broad class of problems that may be partially decomposed into a number of interconnected subproblems. We developed two new classes of methods, both with the unusual characteristic that each subproblem need not be visited at every iteration.

One new method we developed is based on selective linearization and adaptation of ideas from bundle methods. We proved convergence and estimated the rate of convergence of this method. As an additional result, we obtained a rate of convergence estimate for the classical bundle method, thus solving a theoretical problem that had been outstanding for 30 years. These results have been published in the papers:

Y. Du, X. Lin, A. Ruszczynski, “A Selective Linearization method for Multiblock Convex Optimization,” *SIAM Journal on Optimization* 27 (2017) 1102-1117

Y. Du, A. Ruszczynski, “Rate of Convergence of the Bundle Method,” *Journal of Optimization Theory and Applications* 173 (2017) 908-922.

The research was part of the PhD dissertation of Yu Du, defended in May 2017.

We also developed a new method based on block-iterative projective splitting for monotone operators. “Block iterative” means that not all the subsystems being coordinated need be considered at each iteration of the algorithm; instead, calculations may be performed for just a few subsystems, or perhaps just one, between coordination calculations. These algorithms also allow for asynchronous parallel execution. This work is described in the papers:

P.L. Combettes and J. Eckstein, “Asynchronous Block-Iterative Primal-Dual Decomposition Methods for Monotone Inclusions,” *Mathematical Programming* 168(1-2):645-672 (2018).

J. Eckstein, “A Simplified Form of Block-Iterative Operator Splitting, and an Asynchronous Algorithm Resembling the Multi-Block Alternating Direction Method of Multipliers,” *Journal of Optimization Theory and Applications* 173(1):155-182 (2017).

P. R. Johnstone and J. Eckstein, “Convergence Rates for Projective Splitting,” *SIAM Journal on Optimization* 29(3):1931-1957 (2019).

The last of these papers obtains convergence rates in the conventional synchronous case in which all subsystems are considered at each iteration. Four more papers on the general topic of projective splitting are currently in undergoing review or revision for publication.

Application of Traditional Mathematical Programming Methods to Machine Learning Problems

Finally, we initiated a project investigating machine learning application of “classical” mathematical programming methods, originating in the field of operations research, such as column generation and branch-and-bound search. As part of this effort, we devised a new “free form” method for regressing an arbitrary scalar response against a set of observations, obtaining superior results for small datasets. Through column generation, this work combined a linear or quadratic programming master problem with an $\mathcal{NP}$ -hard subproblem called rectangular maximum agreement (RMA). This work is described in the papers:

J. Eckstein, N. Goldberg, and A. Kagawa. “Rule-Enhanced Penalized Regression by Column Generation using Rectangular Maximum Agreement.” *Proceedings of the 34th International Conference on Machine Learning (ICML 2017)*, appearing as *Proceedings of Machine Learning Research* 70:1059-1067 (2017).

J. Eckstein, A. Kagawa, and N. Goldberg. “REPR: Rule-Enhanced Penalized Regression,” *INFORMS Journal on Optimization* 1(2):143-163 (2019).

This work also appears in the PhD dissertation of Ai Kagawa, defended September 2018, along with an application to classification problems and large amounts of material on how to solve the RMA subproblem efficiently in practice. A follow-up paper describing exact and heuristic

algorithms for RMA problems has been delayed by a major service interruption in Rutgers' HPC infrastructure and is currently in preparation.