Kevin J. Ross
e-mail: kjross@stat.stanford.edu
website: http://www-stat.stanford.edu/~kjross
phone: 650-814-0802

Department of Statistics
Stanford University
Sequoia Hall, 390 Serra Mall
Stanford, CA 94305-4065

RESEARCH STATEMENT

1 Introduction.

Stochastic control theory is one active research area which incorporates many of my interests. The basic formulation involves a stochastic dynamical system, often a diffusion process, whose state evolution can be influenced by exercising a control. Associated with the system is a *cost (or reward) function* which may depend on both the state process and control. The main goal is to find the optimal value of the cost, the so-called *value function* (denoted V), and a control policy whose cost achieves this value.

Ideas of stochastic control have found application in a variety of areas such as engineering, economics, communication systems, and mathematical finance. A challenging subclass of problems, *singular stochastic control*, has recently received significant attention because of its role in mathematical finance and stochastic processing systems. Roughly speaking, singular controls are control terms that need not be absolutely continuous (with respect to the Lebesgue measure), and are only required to have paths of bounded variation. Singular controls can be used to model behavior, present in many physical systems, in which applying control can produce large instantaneous changes in system state. (See Section 3.1 for description of a singular control problem in mathematical finance.)

Another common feature of stochastic control problems is *state constraints*, the requirement that the state process take values in some specified set. Singular control problems with state constraints provide a challenging class of problems with many unresolved issues. In this document I summarize some of my research on these problems and in other areas. In Section 2, I describe my work obtaining explicit characterizations for singular control problems associated with controlled stochastic networks. Sections 3 and 4 concern my work in numerical solution methods for singular control problems. In Section 5, I discuss briefly my research contributions in stochastic modeling in biology. Finally, I outline in Section 6 some research I expect to pursue in the future. (Links to [4, 5, 6, 7] can be found at: `http://www-stat.stanford.edu/~kjross/research.html`.)

2 Controlled stochastic networks.

Stochastic networks have been an area of active research in recent years with applications in a variety of disciplines, including manufacturing, communications, and computing. In general, a stochastic network consists of a system in which customers (or jobs, packets, etc.) arrive at random times and are placed in a series of buffers, where they await service by one or more servers. Servers may accept customers from multiple buffers and service completion times may depend on customer characteristics. Upon completion of service, a customer may exit the system or be directed to another buffer to await service there.

A fundamental yet challenging problem with critical practical implications concerns control of multiclass stochastic networks. The main objective is to design a control policy for the network in order to optimize some criteria. Control may include influence over rates of arrival and/or service, and routing or scheduling of jobs. Optimization criteria can incorporate holding costs, server idleness times, and other appropriate performance measures.

Except for simple examples such control problems are quite intractable by classical queueing techniques, and thus suitable approximation methods are needed. One common approach for systems that are critically loaded (see [14]) uses *heavy traffic approximation* to replace, formally, a control problem for a stochastic network by one for a diffusion process. This leads to a challenging class of singular control problems with state constraints. The basic approach is to solve the diffusion control problem and then use insights derived from the solution to obtain a near optimal network control policy.

In [6], Budhiraja and I consider a singular stochastic control problem with state constraints in two-dimensions. In particular, the controlled diffusion process is required to take values in some proper subset $\mathbb{S}$ of $\mathbb{R}^2$. Such a control problem and its connections with queuing networks in heavy traffic has been studied by many authors [15, 24, 3, 9, 25, 19]. The main contribution of [6] is to provide an explicit representation for an optimal control under appropriate conditions on the model parameters.

2.1 Characterization of an optimal singular control.

Explicitly solvable singular control problems are quite rare. In the few examples where explicit solutions are available one finds that an optimal control takes the following form. There is an open set $\mathcal{O}$ in the state space such that starting from within $\bar{\mathcal{O}}$ no control is applied until the state trajectory reaches the boundary $\partial\mathcal{O}$, at which point a minimal amount of push is applied along an appropriate control direction to constrain the state process within $\bar{\mathcal{O}}$. Furthermore, if the initial condition is outside $\bar{\mathcal{O}}$, an instantaneous jump occurs at time 0 that brings the process to $\partial\mathcal{O}$, and subsequently, control is applied as described above. In other words, an optimally controlled process is a reflected diffusion on $\bar{\mathcal{O}}$ with an appropriate (possibly oblique) reflection field. In terms of the associated *Hamilton-Jacobi-Bellman (HJB) equation*, in $\mathcal{O}$ the value function V satisfies a linear elliptic partial differential equation (PDE) and in $\mathcal{O}^c$ a nonlinear first order PDE is satisfied; the boundary $\partial\mathcal{O}$ separating these two regions is referred to as the *free boundary* for the system of PDEs.

Characterizations of optimal singular controls in terms of a diffusion reflected at the free boundary are some of the most useful and elegant results in the field. However, in more than one dimension the only previous such results are due to Shreve and Soner [29, 30]. As one may expect, such characterizations are intimately tied to regularity (i.e. smoothness) properties of the free boundary, which in turn hinge on similar properties of V. In [6], the state constraint feature and lack of regularity of the model parameters (e.g. the *running cost* function ℓ is convex, but not strictly convex nor C^1) make C^2 regularity of V an unrealistic goal. Nevertheless, exploiting the convexity of ℓ we show that V is C^1 in $\mathbb{S}^o$ and its gradient extends continuously to all of $\mathbb{S}$. Our proof is probabilistic and a key ingredient is the availability of an optimal singular control as established in [4].

2.2 Connections with optimal stopping and proof of optimality.

We next turn to the study of the free boundary problem and a representation for an optimally controlled state process. We show that the e_1-directional derivative of the value function is the value function of a closely related *optimal stopping problem*. Connections between singular control problems and optimal stopping were first observed by Bather and Chernoff [2] and subsequently such correspondence results have been studied by several authors [16, 17, 29, 30, 13] in one-dimensional and certain multi-dimensional models. The key differences between previous work and our setting are, again, the presence of both the state constraint feature and the lack of regularity of the model parameters.

The study of the optimal stopping problem suggests an explicit form for the "no action region" $\mathcal{O}$ for an optimal control policy. We establish properties of the corresponding free boundary Ψ; in particular, we show that Ψ is a Lipschitz continuous function. A natural conjecture for an optimally controlled process is a diffusion in $\mathcal{O}$, reflected at $\partial\mathcal{O}$ along an appropriate reflection field. A major obstacle in showing that the conjectured controlled process is optimally controlled is the lack of sufficient smoothness of the free boundary (Ψ is only Lipschitz) and the value function. Typical proofs of such a result (see [29, 30]) follow through an application of Itô's formula using the fact that the value function is a classical solution of the associated HJB equation. In view of unavailability of enough regularity we proceed with a *viscosity solution* approach. The main difficulty in the proof is that due to the lack of sufficient regularity of the value function on $\partial\mathcal{O}$, we cannot apply Itô's formula directly to V. Rather, we consider an approximation V^ϵ of V that is C^2 in an open set containing $\mathcal{O}$, apply Itô's formula to V^ϵ, and finally send $\epsilon \to 0$.

3 Markov chain approximations for singular control problems.

Stochastic control problems can rarely be solved explicitly. Thus, in practice numerical approximations are needed. One approach is to develop numerical schemes based on the associated HJB equation. However, many interesting applications of stochastic control, including problems with singular controls, occur in settings where existence/uniqueness theory of the corresponding HJB equations is not well understood. Even when such theory is available, convergence analysis of the numerical PDE approximation can be quite challenging.

Kushner, Dupuis and co-workers (see [20] and references therein) have developed a powerful machinery for computational problems in stochastic control theory using a probabilistic approach. The main idea is to approximate the original controlled diffusion process by a suitable controlled Markov chain on a finite state space. Next, define an appropriate Markov decision problem (MDP), a discrete time, discrete state analog of the continuous time control problem of interest. "Almost optimal" control policies and value functions for the MDP can be computed using classical iterative procedures such as value space and/or policy space approximations and their refinements. This in turn yields approximations of the value function and optimal control policy for the original control problem. Convergence analysis of the algorithm involves establishing convergence of the value function of the MDP to the value function of the original diffusion control problem as various approximation parameters approach suitable limits. This convergence analysis is completely probabilistic and is based on the theory of weak convergence of probability measures. One main advantage of this approach is that it does not require smoothness of the cost or value functions, nor does it rely on associated HJB equations (which are often difficult to work with due to complex features of the state dynamics).

In [5] Budhiraja and I use the Kushner-Dupuis methodology to develop a convergent numerical scheme for a singular control problem in financial mathematics. While the basic method uses approximation via a suitable Markov chain, specific features of the problem make description of the approximating chain and convergence analysis quite delicate.

3.1 Optimal investment and consumption with proportional transaction costs.

Consider an investor who has two investments available: a risk free asset (bank account) and a risky asset (stock). The investor may instantaneously transfer wealth between the two assets; however, she is charged a transaction fee proportional to the amount of stock bought or sold. She may also consume wealth at some time dependent rate. The investor is endowed with a utility function and her goal is to choose a consumption and portfolio selection strategy which maximizes expected discounted utility over an infinite time horizon. Due to the instantaneous nature of the transactions and their corresponding costs, the stochastic processes which represent cumulative amounts of stock purchases and sales are singular controls. Consumption enters into the model as an absolutely continuous control term. The requirement that the investor must be solvent at all times imposes state constraints.

3.2 Main results.

The problem is quite well studied for certain utility functions (e.g. x^p/p or $\log x$) in which an associated free boundary problem can be solved explicitly (see [12], [31]). For general utility functions, explicit solutions are unavailable and thus numerical approximations are needed. In [5] we prove convergence of a Markov chain based approximation scheme. We also implement the approximation using suitable iterative methods for MDPs, and compute the near optimal consumption control and numerical free boundary of the "no transaction region". There are no concavity, smoothness or growth conditions on the utility function; the key requirement is that the value function is finite and continuous. In particular, uniqueness of the solution of the associated HJB equation is neither needed nor available (in the generality with which the problem is posed). We show that the value functions of a suitable sequence of MDPs converge to the value function of the singular control problem as the approximation parameters approach their limits.

3.3 Discussion of proofs.

There are several challenging aspects of this convergence proof. First, we approximate the original unbounded model by one where the consumption control and state space are bounded. Proving that this approximation yields a consistent scheme is the only place where value function continuity is used. Another substantial difficulty is the state constraint feature of the dynamics, which is particularly problematic since the control directions do not point inward into the state space. Nevertheless, one useful feature of the dynamics is that once the system state reaches the boundary, the only admissible control corresponds to moving the state process instantaneously to the origin and keeping it there forever. This observation allows conversion of the infinite horizon cost to an exit time criterion, which makes some aspects of the convergence analysis simpler. However, degeneracies in state dynamics make treatment of convergence properties of exit times quite subtle.

One major obstacle in proving convergence of value functions for singular control problems is establishing tightness of the sequence of approximating singular control terms in the Skorohod $D[0,\infty)$ space. A powerful technique for bypassing this issue, based on a suitable stretching of time scale, was introduced in [21]. A key ingredient of this technique is to obtain a uniform moment estimate on the approximating sequence of singular controls. In [21] such an estimate follows easily from the form of the cost function, where a strictly positive proportional cost is incurred for exercising singular control. However, in the current problem there is no direct contribution to the cost function, and thus proof of this uniform estimate becomes quite involved. Roughly speaking, the main idea of the proof is that too much use of singular control will push the process to the boundary of the state space.

The final nontrivial part of the analysis is proof of convergence of classical iterative methods to the value function of the approximating MDP. In problems with only absolutely continuous controls, such convergence is established by exploiting a contraction property that is a consequence of the strictly positive discount factor in the cost criterion. However for singular control problems, such contraction estimates are typically unavailable due to the instantaneous nature of the control. To obtain such a contraction estimate, we use once again the special feature of the dynamics, which says that too much use of singular control will rapidly bring the process to the boundary.

4 Efficient computational schemes for finite horizon singular control problems.

Numerical schemes based on Markov chain approximations typically require the solution of a MDP of considerably high dimension, and thus such schemes can be computationally expensive. Several alternative approaches have been presented for singular control problems, most notably [19, 27, 26, 28] in which the free boundary problem is transformed into a sequence of fixed boundary problems. While such methods often offer a computation savings over the corresponding Markov chain approximations, their domain of application is restricted to infinite time horizon problems in which the free boundary does not depend on time. However, many practical applications involve problems with a *finite time horizon.* In such problems the computational complexity is substantially increased since the boundary of the no action region is now a function of the remaining time horizon.

In [22] Lai, Lim, and I exploit connections between singular control and optimal stopping problems to develop a computationally efficient method for solving singular stochastic control problems with finite time horizons. In settings where such a relationship has been established (e.g. [2, 16, 17, 29, 30, 18, 13, 4]), it is observed that the no action region of the original singular stochastic control problem corresponds to the continuation region of an "equivalent" optimal stopping problem, whose value function is the derivative of the value function of the singular control problem with respect to the controlled state. The practical implication of these results is that the computational complexity of singular stochastic control problems, whose solutions require the determination of both *when* to apply control and *how much* control to apply, can be substantially reduced by working with their equivalent optimal stopping problems, whose solutions consist of determining only when to stop.

We begin with a class of singular control problems that can be transformed into optimal stopping problems associated with two-player *Dynkin games*. We introduce a backwards induction algorithm to compute the continuation boundary and the value function for the Dynkin game. We then show how to modify the backwards induction algorithm for more general singular control problems, even in situations in which the singular control problem is not reducible to an equivalent optimal stopping problem because of the presence of additional functions and their derivatives. Convergence of the algorithm is based on corrected random walk approximations to free boundary problems in optimal stopping (e.g. [10, 11, 23]). We illustrate performance of the algorithm in applications in reversible investment, optimal investment and consumption, and controlled stochastic networks.

5 Stochastic modeling in biology.

My work with M.R. Leadbetter, in collaboration with Prof. Karen Burg of the Department of Bioengineering at Clemson University, has introduced me to important issues of stochastic modeling in cell biology (see [7, 8]). The ultimate goal of this research is to utilize techniques of tissue engineering to provide better alternatives toward breast tissue reconstruction for patients who have undergone mastectomy. My contribution toward this goal involved statistical analysis of experimental data and development of realistic models for cell growth, specifically in terms of cell metabolism and lipid production.

6 Future work.

I will continue to research singular stochastic control problems with state constraints. This is a significant class of problems, arising from exciting application areas, that is still in the very early stages of development. I will use insights obtained from two-dimensional singular control problems, and corresponding problems of optimal stopping, in developing techniques for a broader class of networks. The goal here is to characterize the optimally controlled process as a reflected diffusion on a suitable domain associated with a free boundary problem. A key step in this direction is to obtain sufficient regularity of the value function.

Techniques of stochastic modeling and analysis, and stochastic control theory in particular, have already found a strong foothold in many fields such as mathematical finance, computer science, and industrial engineering. Yet, the basic principles are applicable to an even broader array of fields. One of my research goals is to bridge gaps between disciplines to find new areas of application. In many cases, existing theory may be adapted readily to new areas with little more than a change in terminology or interpretation. On the other hand, identifying interesting applications can drive development of innovative probability theory.

Biology is one discipline which holds many opportunities for new applications. One research direction I plan to investigate involves applying queueing theory and diffusion approximation methods to genetic regulatory networks which model interactions of genes, proteins, and RNA and DNA molecules. The first step in this line of research is to identify a suitable queueing network formulation for the biological system of interest. For example, each service facility may represent a specific type of molecule or gene; queue length may represent the number of molecules present or gene expression level; arrivals from outside may represent increases in the number of molecules or expression level due to some external shock; etc. While some work has been done (see [1] and references therein), research is in the very early stages of development. Defining appropriate queueing models for such biological systems will be a significant contribution in itself. Good queueing network models will provide insights, through simulation studies, into the dynamics of genetic regulatory networks. Furthermore, comparing simulation results to microarray data will enhance methods of model selection and validation. Once models are available, techniques of stochastic analysis and queueing theory (diffusion approximations, heavy traffic analysis, etc.) may be employed to model and study behavior of the genetic network. All of the above has potential for exciting and important biological insights.

References

[1] A. Arazi, E. Ben-Jacob, and U. Yechiali. Bridging genetic networks and queueing theory. *Phsyica A*, 332(1): 585–616, 2004.

[2] J. Bather and H. Chernoff. Sequential decisions in the control of a spaceship. *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Vol. III: Physical Sciences*, pp. 181-207. University of California Press, Berkeley, CA. 1967.

[3] A. Budhiraja and A. Ghosh. A large deviations approach to asymptotically optimal control of crisscross network in heavy traffic. *Annals of Applied Probability*, 15(3): 1887–1935, 2005.

[4] A. Budhiraja and K. Ross. Existence of optimal controls for singular control problems with state constraints. *Annals of Applied Probability*, 16(4): 2235–2255, 2006.

[5] A. Budhiraja and K. Ross. Convergent numerical scheme for singular control with state constraints in a portfolio selection problem. *SIAM Journal on Control and Optimization*, 45(6): 2169–2206, 2007.

[6] A. Budhiraja and K. Ross. Optimal stopping and free boundary characterizations for some Brownian control problems. *Annals of Applied Probability*, 18(6): 2367–2391, 2008.

[7] A. Chaubey, K. Ross, M. R. Leadbetter, K. J. L. Burg. Surface patterning: Tool to modulate stem cell differentiation in an adipose system. *Journal of Biomedical Materials Research, Part B: Applied Biomaterials*, 84B(1):70–78, 2008.

[8] A. Chaubey, M. R. Leadbetter, K. Ross, C. Gomillion, K. J. L. Burg. Characterization of the differentiation and leptin secretion profile of adult stem cells on patterned polylactide films. *Journal of Biomaterials Science, Polymer Edition*, To appear.

[9] M. Chen, C. Pandit, and S. Meyn. In search of sensitivity in network optimization. *Queueing Systems*, 44(4): 313–363, 2003.

[10] H. Chernoff. Sequential tests for the mean of a normal distribution. IV. (Discrete case). *Annals of Mathematical Statistics*, 36: 55–68, 1965.

[11] H. Chernoff and A. J. Petkau. An optimal stopping problem for sums of dichotomous random variables. *Annals of Probability*, 4(6): 875–889, 1976.

[12] M. H. A. Davis and A. R. Norman. Portfolio selection with transaction costs. *Mathematics of Operations Research*, 15(4): 676–713, 1990.

[13] X. Guo and P. Tomecek. Connections between singular control and optimal switching. *SIAM Journal on Control and Optimization*, 47(1): 421-443, 2008.

[14] J. M. Harrison. Brownian models of queueing networks with heterogeneous customer population. *Stochastic Differential Systems, Stochastic Control Theory and Their Applications*. Springer, New York, pages 147–186, 1988.

[15] J.M. Harrison and L. Wein. Scheduling networks of queues: Heavy traffic analysis of a simple open network. *Queueing Systems*, 5(4): 265–280, 1989.

[16] I. Karatzas and S. E. Shreve. Connection between optimal stopping and singular stochastic control, Part I: Monotone follower problems. *SIAM Journal on Control and Optimization*, 22(6): 856–877, 1984.

[17] I. Karatzas and S. E. Shreve. Connection between optimal stopping and singular stochastic control, Part II: Reflected follower problems. *SIAM Journal on Control and Optimization*, 23(3): 433–451, 1985.

[18] I. Karatzas and H. Wang. Connections between bounded-variation control and Dynkin games. In *Optimal Control and Partial Differential Equations — Innovations and Applications* (J .L. Menaldi, E. Rofman, and A. Sulem, eds.). Amsterdam: IOS Press, 353–362, 2001.

[19] S. Kumar and K. Muthuraman. A numerical method for solving singular stochastic control problems. *Operations Research*, 52(4): 563–582, 2004.

[20] H. J. Kushner and P. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time.* Springer-Verlag, New York, second edition, 2001.

[21] H. J. Kushner and L. F. Martins. Numerical methods for stochastic singular control problems. *SIAM Journal on Control and Optimization*, 29(6): 1443–1475, 1991.

[22] T. L. Lai, T. W. Lim, and K. Ross. A new approach to singular stochastic control with applications in investment and stochastic networks. Under preparation.

[23] T. L. Lai, Y. Yao, and F. AitSahlia. Corrected random walk approximations to free boundary problems in optimal stopping. *Advances in Applied Probability*, 39(3): 753–775, 2007.

[24] L. F. Martins, S. Shreve, and H. M. Soner. Heavy traffic convergence of a controlled, multi-class queueing system. *SIAM Journal on Control and Optimization*, 34(5): 2133–2171, 1996.

[25] S. Meyn. Sequencing and routing in multiclass queueing networks. II. Workload relaxations. *SIAM Journal on Control and Optimization*, 42(1): 178–217, 2003.

[26] K. Muthuraman. A computational scheme for optimal investment-consumption with proportional transaction costs. *Jounral of Economic Dynamics and Control*, 31(4): 1132–1159, 2007.

[27] K. Muthuraman and S. Kumar. Multidimensional portfolio optimization with proportional transaction costs. *Mathematical Finance*, 16(2): 301–335, 2006.

[28] K. Muthuraman and H. Zha. Simulation-based portfolio optimization for large portfolios with transaction costs. *Mathematical Finance*, 18(1): 115–134, 2008.

[29] S. E. Shreve and H. M. Soner. Regularity of the value function for a two-dimensional singular stochastic control problem. *SIAM Journal on Control and Optimization*, 27(4): 876–917, 1989.

[30] S. E. Shreve and H. M. Soner. A free boundary problem related to singular stochastic control. *Applied stochastic analysis (London, 1989)*, 265–301, Stochastics Monogr., 5, *Gordon and Breach, New York*, 1991.

[31] S. E. Shreve and H. M. Soner. Optimal investment and consumption with transaction costs. *Annals of Applied Probability*, 4(3): 609–692, 1994.