Securities Lending Policy Optimization

Ellis Brown*
Stanford University
elbrown@stanford.edu

Abstract

This paper presents a method to determine an optimal policy for the lending of securities by large institutions in the securities finance market as a final project for the Stanford University AA222 Engineering Design Optimization class. The securities lending process is formulated as a Markov decision process in which the lender decides whether to accept or reject incoming offers from borrowers. This formulation allows for a policy that maximizes the expected return with each decision to be derived using dynamic programming. The framework presented is easily extensible through the creation of more realistic models of the dynamics of the securities lending market.

1 Introduction

The financial world has been taken over by Exchange Traded Funds (ETFs) over the past decade; ETFs currently account for nearly *half* of all trading activity in the United States stock markets [10]. Along the way, a fee war has been incited between major ETF providers vying for market share. With pressure to drive down fees, securities lending—already vital to various financial domains in its own right—has been thrust into a new echelon of importance, as it is one of the primary ways that these providers can offset their operational costs and afford to reduce fees [12]. In order to increase their trading capacity and maximize lending revenue, lenders have increasingly turned to automatic systems to augment their manual operations.

Automatic lending systems today have snowballed into thousands of lines of complex business logic, containing rules that attempt to capture the distilled knowledge of traders and domain experts over the past few decades. As the process of making a lending decision is extremely complex, new edge cases and inefficiencies in the systems materialize regularly and are patched with even more rules. Although these archaic systems are known to be suboptimal, they handle more than half of overall lending volume in some cases. The quality of the decisions that these automatic lending systems make is a limiting factor to the continued growth and competitiveness of major ETF providers; developing a better and more scalable solution is essential to their future prosperity.

In this paper, we propose a framework that allows us to dispense with codified business logic, and solve directly for a policy that maximizes expected return. After imposing a series of simplifying assumptions, we formulate the decision to lend securities as a *Markov decision process* (MDP) that depends on the dynamics of the securities lending market. We can solve such problems using dynamic programming techniques developed by R. Bellman in the 1950s [4]. We present initial results of modeling the distribution of offers from borrowers and solving the formulation using *value iteration*. This framework is intended to be extended with models that relax our simplifying assumptions and better represent reality in order to develop real world policies.

*E. Brown is a part-time graduate student at Stanford University and a full-time Machine Learning Engineer at BlackRock AI Labs. Work email: ellis.brown@blackrock.com.

AA222/CS361: Engineering Design Optimization (Spring 2020), Stanford University, Stanford, CA 94305

2 Background

An ETF is an investment fund—a managed portfolio of securities that investors collectively own—that is traded on an exchange just like a regular security, allowing much greater flexibility to investors than comparable fund types (e.g., mutual funds). ETFs are most often *passively managed*—designed to proportionally track some *index* (essentially a weighted collection of securities, e.g., the S&P 500). The extremely low management fees—due to the minimal costs of passive management, high tax efficiency, and economies of scale—charged by ETFs is their primary advantage to investors beyond the flexibility of intraday trading. As ETF providers must hold the amount of underlying securities that make up the total value of all shares of the fund on the market, they end up with a profusion of holdings. The biggest ETFs currently are each composed of more than $100B worth of securities; the market is dominated by a few giants who manage the majority funds and upwards of $1T each [1, 2, 7]. These giant providers face stiff competition over fees, and have huge incentive to maximally utilize their massive inventories through securities lending in order to offset expenses and maintain competitive pricing.

2.1 Securities Lending

The securities lending process allows owners of securities can earn extra income by lending portions of their holdings out for a fee. The *security* being lent can be any type of financial instrument, such as stocks, bonds, treasuries, or mortgage-backed securities. A securities loan primarily involves two parties: borrowers and lenders. *Borrowers* are typically broker-dealers who need securities in order to perform an advanced investment strategy, such as short selling. *Lenders* are typically large asset-managers or commercial banks who have significant lending operations in place. These operations span seeking out a borrower, negotiating the terms of the loan (such as the fee, duration, and quantity), and facilitating the transfer of securities. Lenders often lend out securities not only on their own behalf, but also on the behalf of third-party investors—such as investment funds, endowments, or insurance companies—as a service. The incentive to lend stems from the guaranteed fee income and controlled risk offered by the lending agreement. [8]

Security loans are typically initiated by the borrower. The lender and borrower negotiate the terms of the loan, including the quantity of securities being borrowed, the lending fee, the duration of the loan, and the nature of the *collateral*. Before receipt of the lent securities, the borrower must send other securities or cash as collateral to the lender; the value of collateral is typically a minimum of 102% of the market value of the securities to be lent [8]. Throughout the duration of the loan, the borrower is free to use the securities however they please—as long as they return the same quantity of the particular security at the end. The borrower typically uses the borrowed securities in order to execute an advanced trading strategy, such as a short. The lender typically reinvests the collateral in low risk investments to earn a small profit. At the end of the duration of the loan, the two parties swap back the securities and collateral, and the borrower pays the lender the predetermined fee.

The fee charged in a securities loan is primarily supply and demand driven, depending on how difficult it is to borrow the security in question [8]. Greater demand for a particular security by borrowers enables lenders to charge a price premium; oversupply by lenders forces them to charge a discounted price. Some securities that are highly available and identically substitutable (e.g., treasuries with the same maturity) earn special designation as *general collateral* (GC). Borrowers are indifferent to which GC security they receive (because they are interchangeable), and so the fee for such assets is a driven by the cost of borrowing cash rather than by supply/demand of the security [3]. The fee attainable by lending GC securities also garners special designation as the *GC rate*. The basket of securities considered GC, along with the GC rate, are determined by market consensus and are in constant gradual flux.

2.2 Automatic Lending

In order make extra money by lending more of their available inventory, lenders have built systems that can automatically lend out securities. Lenders generally automatically accept *all* requests to borrow GC assets at the GC rate determined by the market. This is done via a simple system that plugs in directly to the electronic exchange. As lenders are able to charge significantly higher fees for non-GC assets, almost all effort is allocated to handling such assets appropriately. For the remainder of this paper, it can be assumed that the securities in question are non-GC unless otherwise specified.

Automatic lending systems today primarily use business logic codified by traders and other securities lending domain experts. Many of these systems were created 1–2 decades ago, and have since changed little beyond the steady addition of supplementary rules. These legacy systems are currently used to automatically lend assets on the lower end of the demand/fee spectrum, leaving the most important decisions for the traders to make manually. Though a significant volume of loans are filled by these systems, they are known to be suboptimal and are a huge limiting factor on the efficacy of securities lending operations. An optimized system would not only enable greater revenue on the subset of loans currently being processed by the legacy systems, it would also enable a wider scope of assets (those with higher demand and thus possible fees) to be automatically lent out. This would additionally free up the human traders to allocate more time to the most important offers.

3 Problem Formulation

We first make a series of simplifying assumptions in order to appropriately frame the securities lending problem. We then formulate finding an optimal lending policy as a dynamic programming problem.

3.1 Simplifying Assumptions

We will operate throughout the remainder of this paper from the viewpoint of the lender under the following simplifying assumptions:

- We receive incoming offers to borrow a given security from potential borrowers. The offer contains the desired:
  - *amount* of securities to borrow a
  - *price* (fee) of the loan p
- We have a lendable *inventory* of the given security q
- We can choose to either:
  - *accept* the offer, earning the specified fee but reducing our lendable inventory
  - *reject* the offer, waiting to lend our inventory later at a higher price
- The borrower will hold the borrowed securities for some *duration* d
- Our objective is to maximize the expected future return of lending out securities.

3.2 Distribution Models

We need to model the following distributions:

1. *Offer distribution*: the joint distribution of offers (irrespective of our acceptance/denial)
$$P(p, a) \tag{1}$$
2. *Inventory model*: the conditional distribution of inventory, giving the probability of ending up with quantity q' at the next decision stage given the current quantity q. Changes in inventory can occur from borrowers redeeming outstanding loans and returning the borrowed securities, or for various operational reasons.
$$P(q' \mid q) \tag{2}$$
3. *Duration model*: the conditional distribution of a borrower holding a securities at price p for duration d
$$P(d \mid p, a) \tag{3}$$

To increase the realism of these models, we can include additional features (e.g., an indicator of external market conditions) during the modeling process and then condition on the observed values of the features at decision time.

3.3 Dynamic Programming Problem

We use the stated assumptions and the aforementioned distribution models to formulate the securities lending policy optimization problem as an infinite horizon dynamic programming problem. We define the following:

- The **state** of the system at each decision stage consists of of the incoming offer p, a and our current inventory q. The state $s = \{p, a, q\}$ lies within the *state space* $\mathcal{S}$.
- The **action** act we take at each decision stage lies in the discrete action space $\mathcal{A} = \{\text{accept}, \text{decline}\}$. We can only accept an offer if we have enough inventory $q \geq a$.
- The **discount rate** $\gamma \in [0, 1]$ determines the rate at which our utility derived from future states diminishes. When $\gamma = 1$, we equally value the utility at all future states; as γ decreases from 1, we value the utility at future states less.
- The **expected discounted duration** is the amount of time that we expect for the borrower to hold a securities at price p

$$\bar{D}(p, a) = \sum_{d=0}^{\infty} \gamma^d P(d \mid p, a) \tag{4}$$

- The **action utility function** gives the utility (return) associated with taking action act at state s and then following our policy π. We can only experience a return if we lend. Since we are not able to accept an offer if we do not have enough inventory, we introduce an *invalid action penalty* ρ, which should be some large number (e.g., $1e10$)

$$Q(s, act) = \begin{cases} pa\bar{D}(p, a) - \rho(q < a) & act = \text{accept} \\ 0 & act = \text{decline} \end{cases} \tag{5}$$

- The **state transition function** represents the probability of transitioning to state s' after executing action act at state s

$$T(s' \mid s, act) = \begin{cases} P(q', a')P(q' \mid q - a) & act = \text{accept} \\ P(q', a')P(q' \mid q) & act = \text{decline} \end{cases} \tag{6}$$

- The **policy** $\pi : \mathcal{S} \rightarrow \mathcal{A}$ determines what action act we select at state s.[2]

With these definitions, we can formulate the **expected discounted return** of executing policy π from state s as

$$V^{\pi}(s) = Q(s, \pi(s)) + \gamma \sum_{s'} T(s' \mid s, \pi(s)) V^{\pi}(s') \tag{7}$$

We want to find the **optimal policy** π^*, such that the expected discounted return is maximized:

$$\pi^*(s) = \arg\max_{\pi} V^{\pi}(s) \tag{8}$$

4 Method

In this section we show initial distribution modeling results and explain our methodology for finding an optimal policy.

4.1 Distribution Modeling

In order for our optimized policy to be successful in reality, we need to create realistic models of the distributions defined in Section 3.2. We expect this to require significant effort, and so again start with a simple approach and leave improvements to future iterations. We show initial results in modeling the offer distribution (1) in Section 4.1.1. Due to time constraints, we use simple toy models for

[2]Note: we are using a stationary policy; future iterations could explore policies that take into account the past history of states $s_{0:t}$ and actions $a_{0:t}$.

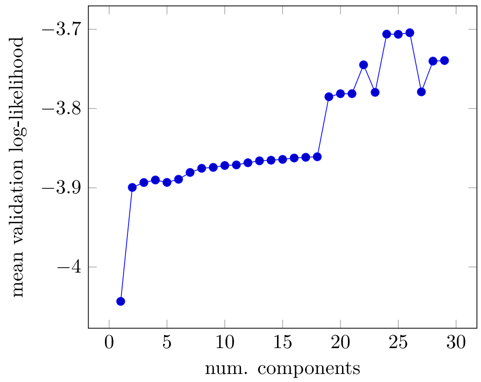

Figure 1: We use 5-fold cross validation to select GMM hyperparameters (number of components) that maximize the log-likelihood. The log-likelihood here is the mean across all folds, evaluated on the validation set. This search results in $n = 26$ components.

the inventory (2) and duration (3) models. For the inventory model, we use a truncated[3] Normal distribution centered on the current quantity, with a standard deviation $\sigma = 1$. For the expected discounted duration, we use a simple linear function of price and amount $\bar{D}(p, a) = 5 + 0.05a - 0.75p$. We intentionally keep the expected duration low for the sake of easy simulation, real world durations would be significantly longer.

4.1.1 Offer distribution

We obtained real offer data from BlackRock's Securities Lending desk. The offer data includes the price and quantity of shares requested by the borrower. The distribution of each variable is highly right-skewed, and so we apply a log transformation to each. We then remove offers that contain an outlier ($|Z| > 3$) on either value in order to reduce noise. There are explanations for each of the outliers, and they will need to be addressed by more realistic models.

We fit a *Gaussian mixture model* (GMM) to our preprocessed data using the *Expectation-Maximization* (EM) algorithm. We use *k-fold cross validation* with 5 folds to select the hyperparameter of our GMM (see Figure 1) [6]. We evaluate our best model ($n = 26$ components) on a held out test set, getting a log-likelihood of -3.67. Figure 2 shows a visualization of the probability distribution obtained.

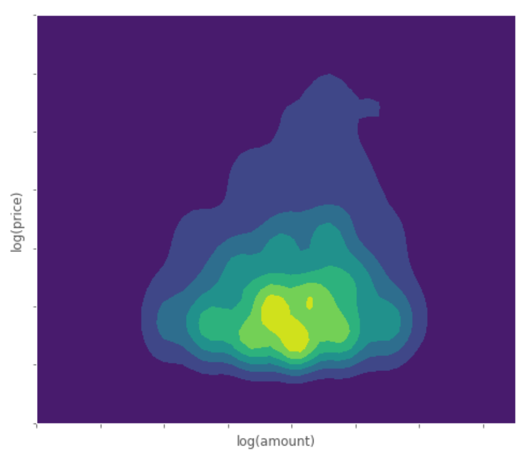

(a) Contour

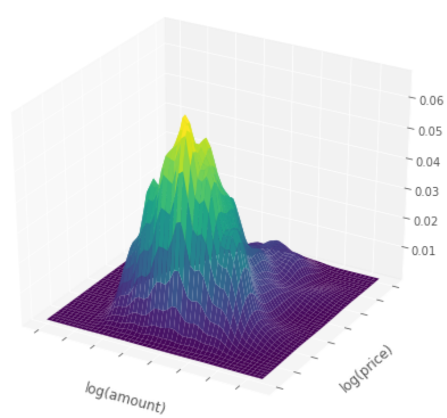

(b) Surface

Figure 2: The joint probability distribution over offers $P(p, a)$, fit to real data using a GMM with $n = 26$ components, selected using cross-validation. *Due to the confidential nature of BlackRock's data, all value markers are omitted.*

[3]We truncate between 0 and ∞ in order to enforce nonnegativity, as we cannot have negative inventory.

4.2 Baseline Stochastic Policy

Before developing an intelligent policy, we first create the simplest possible policy in order to serve as a benchmark. We define a *stochastic policy* that accept all offers with some fixed probability $\phi \in [0, 1]$, assuming there is adequate inventory

$$\pi_{\text{stoch}}(s) = \begin{cases} \text{accept} & q > a \text{ and } \phi > r \\ \text{decline} & \text{otherwise} \end{cases} \tag{9}$$

for some uniformly distributed random number $r \sim U(0, 1)$. Note that stochastic policy makes no regard to the price of the offer, and as such we would expect it to perform very poorly.

4.3 Value Iteration

We want to develop an optimal policy (12). According to Bellman's *principle of optimality*, "an optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision" [4]. For infinite horizon problems, like the one in question, this implies that if we can estimate the optimal expected discounted return (7) of all states attainable from the current state, the action that leads to the state with maximal return will be the optimal action. Formally, an optimal policy can be proven to satisfy the *Bellman equation* [5]

$$V^*(s) = \max_{act} \left(Q(s, act) + \gamma \sum_{s'} T(s' \mid s, act) V^*(s') \right) \tag{10}$$

Note that V^* is a recursion. We can approximate this value using *value iteration*, which is essentially just iterative application of (10)

$$V_{k+1}(s) = \max_{act} \left(Q(s, act) + \gamma \sum_{s'} T(s' \mid s, act) V_k(s') \right) \tag{11}$$

until V converges. Value iteration is proven to converge with for any bounded initialization (i.e., $|V_0(s)| < \infty, \ \forall \ s \in \mathcal{S}$) [5]. As our intuition implied, armed with an estimate of V^*, the optimal policy is to simply choose the action that leads to the maximal expected discounted return:

$$\pi^*(s) \leftarrow \arg\max_{act} \left(Q(s, act) + \gamma \sum_{s'} T(s' \mid s, act) V^*(s') \right) \tag{12}$$

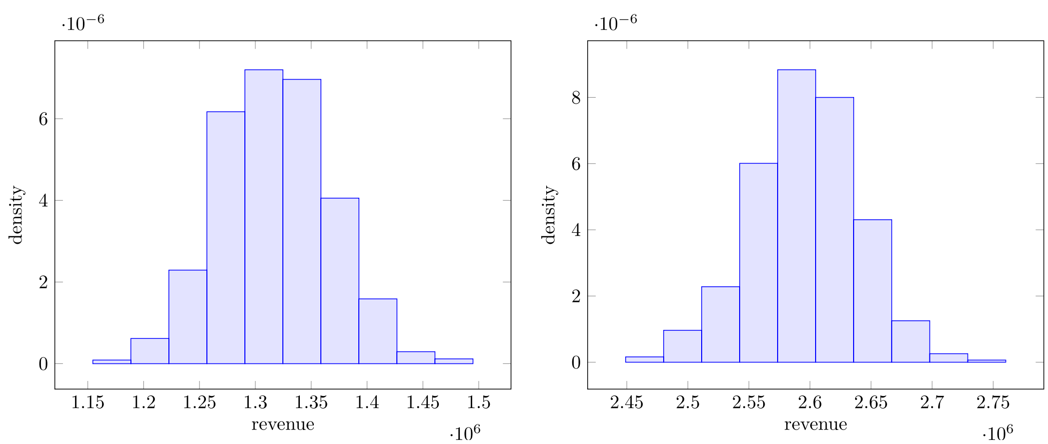

(a) Stochastic Policy π_{stoch}

(b) Value Iteration Policy π^*

Figure 3: The distribution of revenue received over 1000 simulations of 1000 decision stages per policy. The stochastic policy is executed with $\phi = 0.5$, making the decision to lend a coin flip. The mean revenue using our optimal policy is nearly double that of the stochastic policy.

5 Initial Results

In this section, we present initial results comparing the baseline stochastic policy with our optimal policy found through value iteration. In order to visualize the policies in action, we run simulations with initial inventory $q = 35$, and display the state and rewards over 50 decision stages in Figure 4 (next page). We can see that the optimal policy has learned to decline lowball offers, for example at decision step 2,5, and 8. This allows it to better utilize the available inventory and earn more revenue.

We compare the revenue of both policies over 1000 simulations in Figure 3. We can see that the optimal policy learned through value iteration is able to bring in nearly double the revenue compared to the baseline stochastic policy.

6 Conclusion

The problem of lending securities was formulated as a Markov decision process, in which the lender decides whether to accept or reject incoming offers from borrowers. This formulation provides a rich framework upon which we can devise policies that maximize the lenders return through dynamic programming. It is easily extensible through the creation of more realistic models of the dynamics of the securities lending market. Initial results of fitting an model of incoming offers to real world data were shown. Finally, we compared a policy optimized through value iteration to a baseline stochastic policy in simulation.

Acknowledgments

This project is joint work with an early stage project at BlackRock AI Labs that I am working on alongside one co-worker, Mladen Laudanovic; we are being advised by Mykel Kochenderfer, Rob Tibshirani, and Stephen Boyd. Mykel contributed an initial formulation of the problem that was adapted in this paper, along with an amazing framework to generate and evaluate policies. As the success of the project will hinge on realistic models of the distributions in Section 3.2, I have focused the lion's share of my time on understanding the domain as well initial attempts at modeling the distributions.

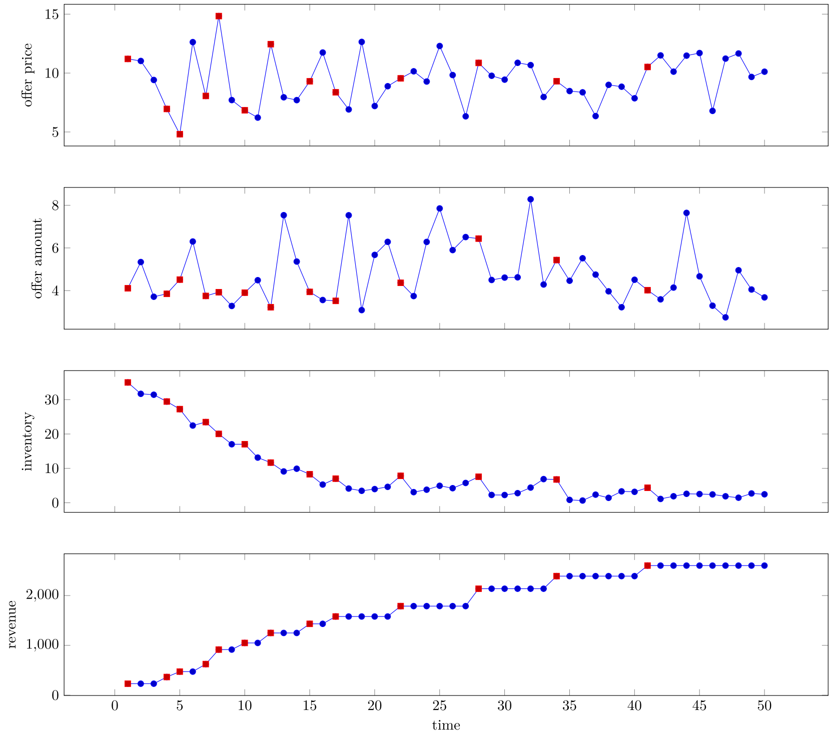

(a) Stochastic Policy π_{stoch}

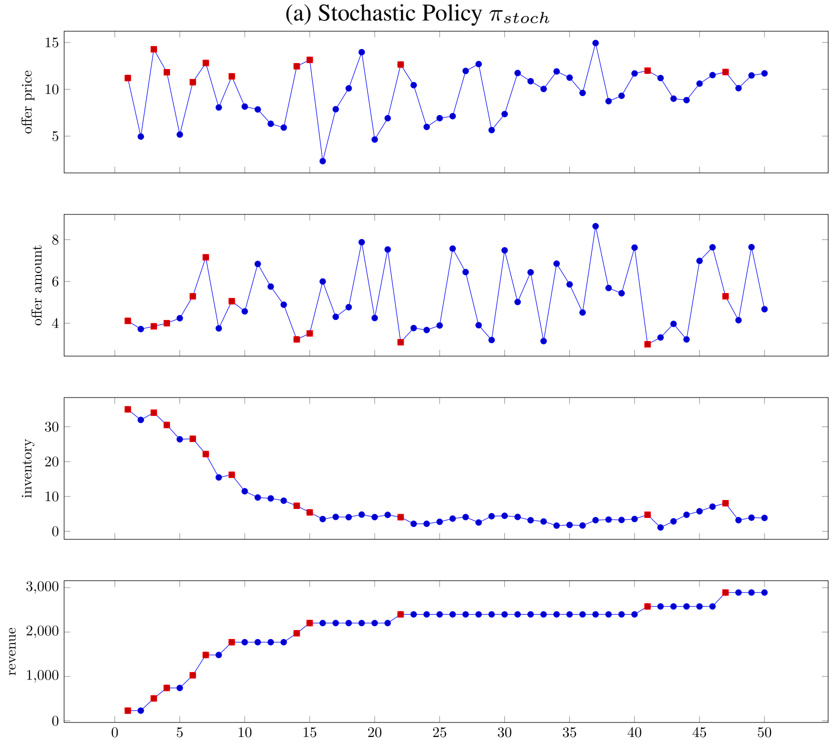

(b) Value Iteration Policy π^*

Figure 4: Example simulation of both policies. Each plot displays the state (price, amount, quantity) and revenue at each decision stage. Initial inventory $q = 35$. Red dots correspond to accepted offers.

References

[1] ETF issuer league tables. https://etfdb.com/etfs/issuers/. Accessed: 2020-6-6.

[2] Largest ETFs: Top 100 ETFs by assets. https://etfdb.com/compare/market-cap/. Accessed: 2020-6-6.

[3] What is general collateral (GC)? https://www.icmagroup.org/Regulatory-Policy-and-Market-Practice/repo-and-collateral-markets/icma-ercc-publications/frequently-asked-questions-on-repo/8-what-is-general-collateral-gc/. Accessed: 2020-6-6.

[4] Richard Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.

[5] Mykel J Kochenderfer. *Decision Making Under Uncertainty: Theory and Application*. MIT Press, 2015.

[6] Mykel J Kochenderfer and Tim A Wheeler. *Algorithms for Optimization*. MIT Press, March 2019.

[7] Mark Kolakowski. Who are the ETF giants? https://www.investopedia.com/who-are-the-etf-giants-4691723, July 2019. Accessed: 2020-6-6.

[8] Paul Lipson, Bradley Sabel, and Frank M Keane. Securities lending. Technical Report 555, Federal Reserve Bank of New York, March 2012.

[9] David Loader. Chapter 7 - securities lending. In David Loader, editor, *Clearing, Settlement and Custody (Third Edition)*, pages 159–175. Butterworth-Heinemann, January 2020.

[10] Chris Newlands and John Authers. Exchange traded funds: taking over the markets. *Financial Times*, December 2016.

[11] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. Wiley, New York, 1994.

[12] Nathan Reiff. The role of securities lending in ETF returns. https://www.investopedia.com/investing/role-securities-lending-etf-returns/, April 2018. Accessed: 2020-6-5.