



Yale ICF Working Paper No. 00-64

October 2001

DYNAMIC TRADING POLICIES WITH PRICE IMPACT

Hua He

Lehman Brothers Inc. Tokyo, Japan

Yale School of Management

Harry Mamaysky

Yale School of Management

This paper can be downloaded without charge from the
Social Science Research Network Electronic Paper Collection:

<http://papers.ssrn.com/abstract=289257>

DYNAMIC TRADING POLICIES WITH PRICE IMPACT*

HUA HE †
Lehman Brothers Inc.
Tokyo, Japan
and
Yale School of Management
New Haven, CT

and

HARRY MAMAYSKY ‡
Yale School of Management
New Haven, CT

First version: July 2001
Last revised: October 2001

*We would like to thank Werner Stanzl for helpful comments and discussions.

†hua.he@yale.edu

‡harry.mamaysky@yale.edu

DYNAMIC TRADING POLICIES WITH PRICE IMPACT

Abstract

In this paper we analyze the optimal policy for a risk averse agent who wants to sell a large block of shares of a risky security in the presence of price impact and transactions costs. Our framework reduces to the standard Merton portfolio problem in the absence of any market frictions. Optimal liquidation results in revenue distributions which are substantially different from those generated by a naive strategy. The main tradeoff involves choosing between revenue distributions which have high means versus those which have low variances. Furthermore, our results suggest that the effective liquidity of a security depends on its return distribution and on the characteristics of the agent carrying out the trade, as well as on the price impact function.

JEL Classification: G11, G12.

1 Introduction

The study of optimal trading policies in securities markets with continuous trading has been one of the central topics in the field of financial economics. Merton's pioneering work (1969, 1971) has laid the groundwork on this topic. An important result in Merton (1971) is that investors optimally invest a constant proportion of wealth in risky assets when the asset returns are lognormally distributed and the utility function exhibits constant relative risk aversion. This result provides valuable insight and theoretical guidance as to how investors should dynamically allocate funds across different assets or asset classes. Merton's work on dynamic trading has made a significant impact on the investment management profession.

One of the key assumptions in Merton's model is that investors are price-takers and they trade competitively. As such, their dynamic trades – regardless of the size – do not affect the prices of the traded securities. While the assumption of no price impact may not be practically important for investors making allocation decision over a very long time horizon, price impact can make a significant difference when investors need to execute large trades over a short time horizon. With price impact, trades affect the current execution price as well as the future execution price if such impact on price is permanent. Consequently, price impact can affect investors' optimal trading behavior significantly. The market microstructure literature has shown both theoretically and empirically that large trades move the price of the underlying securities, either for informational or liquidity reasons.¹ While optimal trading policies in a Merton-type economy are well understood, to date little is known about how investors' dynamic trading policies should be constructed to take into account the price impact associated with large trades.

In this paper, we study an individual's optimal trading problem in a continuous-time partial equilibrium setting in which the prices of the underlying securities are affected by his trades. Specifically, we assume that investors face a upward sloping supply curve such that prices are pushed up as they buy shares and moved down as they sell shares. In a setting similar to that of Merton (1971), we analyze an investor's optimal trading problem when he faces a linear supply curve. While Merton's model focuses primarily on how an investor optimally allocates funds across different assets over time, we address the issue of how to optimally implement a large trade, given an allocation decision. Specifically, we consider the problem of how to optimally execute a block trade: Our investor needs to acquire or unwind a large block of stocks over a short period of time. Using the technique of stochastic control, we provide a complete solution to the block-trading problem specified here.

The optimal block-trading problem described above has a wide range of practical applications. First, it can be used for implementing a portfolio re-balancing strategy at a market microstructure level. Portfolio managers frequently face the need to alter their portfolio allocation mix as their estimates of expected returns and risks are revised. As a result, they frequently are faced with the task of buying or selling a large block of a single stock or multiple stocks in a short time-period. The

¹The classic theoretical study of price impact due to informational asymmetry is Kyle (1985). See also Vayanos (2000). Empirical studies of price impact include Holthausen, Leftwich, and Mayers (1990), Keim and Madhavan (1996), and Chan and Lakonishok (1995).

optimal block-trading policies developed in this paper provide a solution for implementing their allocation decision at a micro level. Second, the liquidity discount implied by our model provides theoretical guidance for broker dealers to quote the price of a principal trade. In a principal trade, a customer would trade a large block of a single stock or multiple stocks with a broker dealer. The dealer is asked to bid (or offer) for a large block of stock(s) from a customer (often at the close of the day) at a discount (or premium) from the prevailing market price. The implicit price discount (or premium) suggested by our model offers dealers a theoretical tool for deriving a fair discount or premium.² Furthermore, our model shows how this discount or premium reflects a given broker-dealer's risk preferences, which may reasonably be expected to vary greatly from one broker-dealer to the next.

In order to capture the idea that a broker-dealer would like to sell a large block of shares, we proceed as follows: Agents in our model start off holding shares of stock and some amount of cash. Share prices evolve with uncertainty over time, while cash appreciates at some riskless interest rate. We endow agents with a utility function over their cash holdings at the end of some prespecified time period. However, we assume that at the end of this time period, an agent derives no utility from any shares of stock which remain in his portfolio. For example, to induce an agent to sell off his share holdings over a one year time period, we assume that the agent cares only about his cash holdings one year from now. Since the agent is indifferent as to how many share of the stock he holds in one year, he has an incentive to liquidate his initial block of shares over the course of the next year. Any time that an agent sells some number of shares, the revenue from that transaction is placed into that agent's cash account, and grows at the riskless rate. In this paper, we study such an agent's optimal liquidation policy.

For this liquidation policy to be nontrivial, we need to create an incentive for agents to split up their trades over time. One such incentive arises from the fact that the amount of price impact incurred from a given sale varies over time, reflecting the ebbs and flows of liquidity in the market. Another reason to split up trades has to do with the possibility that selling a share today and a share tomorrow may incur less price impact (but more uncertainty) than selling both shares today. While our setup can accommodate both of these effects, our numerical analysis focuses on the latter.

In order to understand the properties of a given liquidation strategy, it is useful to characterize that strategy by the distribution of cash holdings which it produces at the end of the liquidation period. In other words, for a single path of the economy, an agent's trading strategy will produce some amount on money in one year's time. Different paths produce different realizations of this money amount. Therefore, we can describe a liquidation strategy by the distribution of cash revenues which it produces. Our analysis suggests that agents must decide between liquidation strategies which produce revenue distributions with high means versus those which produce revenue distributions with low variances. Hence the classic risk-return tradeoff manifests itself in our analysis. Since (by assumption) the expected return on stocks is higher than the riskless rate, high mean strategies require agents to hold off selling shares for as long as possible. On the other hand,

²Dealers making active markets in principal trades can improve their prices over our theoretical guideline if they can cross between buy and sell orders.

low volatility strategies require heavy initial selling in order to reduce risk, but then result in lower expected revenues.

Since we solve for the value function of an optimizing agent who wants to sell a block of shares, it is rather straightforward to compute the shadow price (of the agent) for those shares. That is, if an agent has no dollars and one thousand shares, how many dollars (say C) would it take before he was willing to hold no shares, but C dollars. Clearly this amount captures the expectation and risk of the optimal liquidation policy which that agent would have followed to sell the thousand shares. We can think of $C/1000$ as the shadow price for each share of the block. This price, typically lower than the current market price, therefore provides a measure of the illiquidity of the market for a given security, as a function of the security characteristics, the price impact function, and the risk preferences of the agent in question. We find that the slope of the price impact function, which is the typical measure of market liquidity used in the finance literature, is actually not a very good measure of liquidity. It is possible that an agent would be willing to demand a higher per share price to part with a stock with a higher price impact than for a stock with a lower price impact, depending on the risk-return characteristics of those two securities. Furthermore, it is possible that stock A fetches a higher price than stock B when the agent holds one hundred shares, but that this relationship reverses once the share holdings increase. Our analysis suggests that the drift and volatility of a security, the number of shares to be sold, as well as the risk preferences of the agent who will carry out the trading, are important determinants of the effective liquidity of that security. Knowing only the slope of the price impact function, therefore, is not sufficient.

There are several important papers in the literature that are closely related to the work presented here. Bertsimas and Lo (1998) study an optimal trade execution problem similar to us in a discrete time setting. Their objective is to maximize the expected revenue (i.e., risk neutral preferences) assuming a linear price impact function, and they derive optimal trading policies using stochastic dynamic programming. In the case when the asset price process follows a zero-drift random walk a linear price impact function, they show that the optimal execution strategy is to break the trade equally across time. As shown in this paper, this strategy is suboptimal if preferences are not risk neutral. Stanzl and Huberman (2001) extend the analysis of Bertsimas and Lo to penalize liquidation strategies which are volatile. While their analysis takes into account the idea that risky strategies are undesirable, they do not explicitly model investor preferences, and hence cannot solve for liquidation strategies which are consistent with utility maximization. Almgren and Chriss (1998) study a similar problem in a discrete time setting with both permanent and temporary market impact. They characterize the mean and variance of the revenue function, and consider optimal trading policies in the form of a mean-variance efficient frontier associated with the revenue function. In an important paper, Subramanian and Jarrow (2001) present a model of liquidity discount in a continuous time setting that is quite similar to ours. In their model, an investor is also faced with the problem of selling a large block of stock. They introduce a notion of execution time to ensure that investors trade at discrete time intervals. Since they are primarily concerned with the expected total revenue generated from trades, risk neutral preferences are effectively being

imposed so that money market investment does not play a role.³ In contrast, the decision rules in our model with risk aversion always involve a tradeoff between investment in the risky assets and investment in the riskless asset. Our setup allows for a direct comparison with Merton's model with no liquidation constraints or price impact.⁴

The rest of this paper is organized as follows. Section 2 sets up the model. Section 3 characterizes the optimal solution to the investor's control problem. Section 4 presents a closed form solution for the case of risk-neutral preferences with continuous trading. Section 5 specializes the characterization of the solution to the case of an investor with a power utility function. Section 6 shows how our formulation of the investor's problem with price impact reduces to the standard Merton formulation with no price impact. Section 7 analyzes the solution of the general case of our model. Finally, Section 8 suggests how our work can be extended.

2 The Model

Consider a dynamic trading problem in which a risk averse agent wants to sell a block of \bar{X} shares of a risky stock over a fixed time interval $[0, T]$. The agent has the option of selling all \bar{X} shares at time 0 or at time T , or alternatively, he can break the block trade into small slices and sell them over time. Since continuous trading is allowed, the agent can trade both continuously and discretely. The agent's problem is to sell his holding in such way that it optimizes the expected utility of the total revenue generated from the transactions.

(Revenue Function) We begin by defining a revenue function associated with the agent selling δ shares of the stock at the market (e.g., via a market order). In the case of competitive trading with no price impact, if the current price of the stock is P , then the revenue generated from the sale would simply be $\delta \times P$. With the presence of price impact, the revenue generated from this transaction will be less than $\delta \times P$. As a notation, we define a revenue function $R(\delta, P)$ which indicates how much money an agent is able to obtain from an immediate sale of δ shares when the pre-trade price is P . An example of a revenue function which reflects price impact is $R(\delta, P) = \delta \times P e^{-\lambda \delta}$, for some constant λ . Hence selling more shares lowers the per share price for all shares in the order. The revenue function therefore reflects, among other things, the price impact function (here $e^{-\lambda \delta}$). We will introduce other forms of the revenue function later.

(Money Market Account) We assume that there are two assets available for investing. The first asset is the risky stock whose price dynamics will be discussed below. The second asset is a riskless bond, yielding a constant rate of interest r . This asset is essentially a money market instrument. The agent can always invest his revenue from stock sales into the money market account to earn interests. We use $M(t)$ to denote the agent's holding of the money market account at time t . It will become clear below why introducing the money market account is an important step for trade decisions in the presence of price impact.

³Their extension to the case with risk aversion requires the investor to consume the revenue from sales immediately after the transaction.

⁴Another related paper, Liu and Yong (2001), analyzes the pricing of options in a setting similar to our own.

(Stock Holdings) Since we are studying the problem of an agent trying to liquidate a large block of shares, we will only consider policies in which the agent is allowed to make sales. Buying additional shares is not allowed. The agent's stock holdings as of time t are given by

$$X(t) = X(0) - S(t) - \sum_{\{k:\tau_k \leq t\}} \delta_k, \quad (1)$$

where $X(0) = \bar{X}$, $S(0) = 0$, $S(t)$ is an increasing and continuous process, and $\{\tau_k\}$ are a collection of stopping times at which the agent sells a discrete number of shares $\delta_k > 0$. In other words, the agent can trim his holdings both continuously and discretely.

(Price Impact) The time t price of the risky stock is given by

$$P(t) = P(0) + \int_0^t [\mu P(s)ds + \sigma P(s)dw(s) - \lambda P(s)dS(s)] + \sum_{\{k:\tau_k \leq t\}} \Delta P(\delta_k, P(\tau_k)), \quad (2)$$

where λ is a constant⁵ and $\Delta P(\tau_k, P(\tau_k))$ is the change in the stock price which occurs when the agent trades a non-infinitesimal amount. Without discrete trading, the dynamics of the price process can be represented by the stochastic differential equation:

$$dP(t) = \mu P(t)dt + \sigma P(t)dw(t) - \lambda P(t)dS$$

This implies that the time t price is given by

$$P(t) = P(0) \exp\left[\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma w(t) - \lambda S(t)\right], \quad (3)$$

where $S(t)$ is the cumulative sales through time t . In particular, for any sales policy $dS(t)$ which satisfies $\int_0^T dS(t) = S(T)$, the time T price will be identical.

The price impact due to continuous sales is linear in the dollar amount sold. For discrete trading, we assume the following price impact function:

$$\Delta P(\delta, P) = P \times e^{-\lambda\delta} - P. \quad (4)$$

However, as will become apparent, our approach can handle much more general price impact functions in the discrete trading case.

(Wealth Dynamics) We track the agent's wealth by tracking his money market account. This is given by the following equation:

$$M(t) = M(0) + \int_0^t rM(s)ds + \int_0^t (P(s) + dP(s))dS(s) + \sum_{\{k:\tau_k \leq t\}} R(\delta_k, P(\tau_k)), \quad (5)$$

Note that the sale of dS shares occurs at the post-impact price of $P + dP$. However, because we have restricted $S(t)$ to be increasing, and therefore of finite variation, the term $dPdS = 0$. In contrast to the Merton-type problem when $W(t) = M(t) + X(t) \times P(t)$ is often used to track the

⁵We take the price impact parameter as given in this paper.

agent's wealth, we have to track $M(t)$ directly due to the presence of price impact as the agent will not be able to realize the value of the stock holding, as he sells his holding in the market. This suggests that $M(t)$, $X(t)$ and $P(t)$ are state variables that are necessary to characterize the agent's optimal trading problem.

(Risk Aversion Utility Function) As assumed above, the agent is interested in liquidating \bar{X} shares of a risky security over the time interval $[0, T]$. We therefore define the agent's preferences as the expected utility over his time T wealth or revenue accumulated over the course of liquidation,

$$\mathbf{E} \left[u \left(M(T) + R(\delta, P(T)) \right) \right] \quad (6)$$

where we require that $\delta \leq X(T)$, i.e., we allow free disposal. The utility function u is assumed to be monotonically increasing and concave. The investor's problem is to maximize his expected utility of consumption of money at the final date subject to the stock holding constraint (1), the price impact function (2), and the money market account constraint (5).

We will show later that for the preference class with constant relative risk aversion we are able to reduce the number of state variables to simplify the optimization problem significantly.

(Indirect Utility) The three state variables M, P, X , and time t fully describe the state of the world for the agent at time t . The agent's time t value function is given by

$$J(M, X, P, t) = \sup_{S, \{\tau_k, \delta_k\}} \mathbf{E}_t \left[u \left(M(T) + R(\delta, P(T)) \right) \right], \quad (7)$$

where the dynamics of M, X and P are as defined above and \mathbf{E}_t is the expectation operator with respect to time t information.

(Boundary Conditions) It is easy to see that the boundary condition at time T is given by

$$J(M, X, P, T) = \max_{0 \leq \delta \leq X} u \left(M + R(\delta, P) \right). \quad (8)$$

Since only sales are allowed, having zero shares of the risky asset implies that at time T , only the cash account will be non-zero. Hence we have

$$J(M, 0, P, t) = u \left(M e^{r(T-t)} \right). \quad (9)$$

From (2), we see that $P = 0$ is an absorbing barrier, which renders the holdings of the risky asset worthless, and gives us the following boundary condition

$$J(M, X, 0, t) = u \left(M e^{r(T-t)} \right). \quad (10)$$

3 Characterizations and Optimality Conditions

In this section we characterize the optimal solution of the dynamic trading/liquidation problem specified in the previous section. We consider two cases. First, we assume that transactions are costless in the sense that there are no fixed costs associated with sales. In this case, trades can be executed both continuously and discretely. In the second case, we introduce fixed transaction costs which induces discrete trading.

3.1 The Continuous Trading Case

Let us assume that transactions are costless. Selling shares incurs price impact, but is costless outside of that. In particular this implies that $R(0, P) = \lim_{\delta \downarrow 0} R(\delta, P) = 0$, i.e. the revenue function is right-continuous at 0. This section will provide a heuristic derivation of necessary conditions for the optimal policy. We will provide a formal proof in the Appendix that the conditions set out in this section are sufficient for optimality.

The state space is divided into three regions. The first is the no-trade region, meaning that when the state variables are inside this region the agent will optimally take no action. The second region is the trade region, and if the state variables ever enter this region the agent will trade to bring the state variables back into the no-trade region. The third region is the boundary between the first two.

The value function will satisfy a Bellman-type equation in the no-trade region. For the differential operator \mathcal{L} given by

$$\mathcal{L}[J] = \mu P J_P + \frac{1}{2} \sigma^2 P^2 J_{PP} + r M J_M, \quad (11)$$

the Bellman-type equation is

$$J_t + \mathcal{L}[J] = 0. \quad (12)$$

Here subscripts denote partial derivatives. At any point where agents decide to trade it must be that the value function immediately prior to the trade is equal to the value function immediately after the trade. Formally this implies that

$$J(M, X, P, t) = J(M + R(\delta, P), X - \delta, P + \Delta P(\delta, P), t), \quad (13)$$

where δ is the trade amount and $\Delta P(\delta, P)$ is the price impact function. Note that we have assumed that $\Delta P(\delta, P) = P \times e^{-\lambda \delta} - P$. This form of price impact is consistent with (2) if an agent were to sell δ shares in a continuous fashion. Optimality of the trade amount implies that the following first-order condition must be satisfied

$$0 = J_M(M + R, X - \delta, P + \Delta P, t) R_\delta(\delta, P) - J_X(M + R, X - \delta, P + \Delta P, t) - \lambda P e^{-\lambda \delta} J_P(M + R, X - \delta, P + \Delta P, t). \quad (14)$$

For an infinitesimal trade (i.e. $\delta \approx 0$) to be optimal, we must have that

$$0 = J_M(M, X, P, t) R_\delta(0, P) - J_X(M, X, P, t) - \lambda P J_P(M, X, P, t). \quad (15)$$

Note that in (14), (M, X, P, t) is assumed to be in the trade region while $(M + R, X - \delta, P + \Delta P, t)$ is at the boundary. In (15), (M, X, P, t) is assumed to be at the boundary. These two equations suggest that

$$R_\delta(\delta, P) = R_\delta(0, P + \Delta P) \quad (16)$$

This equation imposes a restriction on the function form of the revenue function R . Indeed, if we limit our search for separable revenue functions of the form $R(\delta, P) = P f(\delta)$, then it is easily shown that the solution to (16) for separable revenue functions is given by

$$R(\delta, P) = \frac{P}{\lambda} \left(1 - C e^{-\lambda \delta} \right), \quad (17)$$

for some constant C . For $C = 1$ this corresponds to the revenue from selling each of δ shares at its marginal price (i.e. the n^{th} share, with $n < \delta$, is sold at a price which reflects the sale of only n , and not δ , shares). That is

$$\frac{P}{\lambda} (1 - e^{-\lambda\delta}) = \int_0^\delta P e^{-\lambda n} dn$$

We can consider the revenue function in (17) with $C = 1$ as the revenue generated from the good liquidation.

A heuristic argument suggests that condition (15) should hold in the no-trade region as well. Noticing that the optimal δ is a function of M, X, P, t , and that $\partial J / \partial \delta = 0$. Hence we have that

$$\begin{aligned} J_M(M, X, P, t) &= J_M(M + R, X - \delta, P e^{-\lambda\delta}, t), \\ J_X(M, X, P, t) &= J_X(M + R, X - \delta, P e^{-\lambda\delta}, t), \\ J_P(M, X, P, t) &= J_M(M + R, X - \delta, P e^{-\lambda\delta}, t) R_P(\delta, P) \\ &\quad + e^{-\lambda\delta} J_P(M + R, X - \delta, P e^{-\lambda\delta}, t). \end{aligned}$$

Plugging these into equation (14) and using (17), it is easy to verify that (15) holds for (M, X, P, t) in the no-trade region as well.

Combining (12) and (15), we conclude that in general the value function satisfies the following differential inequality

$$\max\left(J_t + \mathcal{L}[J], J_M P - J_X - \lambda P J_P\right) = 0 \quad (18)$$

This means that the state space will be divided into three distinct regions as follows:

Non-Infinitesimal Trade:

$$J_t + \mathcal{L}[J] < 0, \quad J_M P - J_X - \lambda P J_P = 0;$$

Infinitesimal Trade:

$$J_t + \mathcal{L}[J] = 0, \quad J_M P - J_X - \lambda P J_P = 0;$$

No Trade:

$$J_t + \mathcal{L}[J] = 0, \quad J_M P - J_X - \lambda P J_P < 0.$$

We summarize our results in the following theorem for sufficiency, the proof of which is deferred to the Appendix.

Theorem 1 *Let $J(M, X, P, t) \in C^{2,2,2,1}$ be a solution to the differential inequality (18) with the boundary conditions (8-10). If J satisfies a set of regularity conditions, then $J(M, X, P, t)$ is the indirect utility function defined as in (7). Moreover, the optimal policy for X is to do nothing in the no-trade region and trade to the boundary in the trade region.*

3.2 The Case of Discrete Trading with Frictions

Let us now assume that every time a sale is initiated, a cost of $\kappa \geq 0$ units of money must be paid. We first note that not trading incurs no cost, or that $R(0, P) = 0$. An implication of these

assumptions is that $R(0, P) - \lim_{\delta \downarrow 0} R(\delta, P) = \kappa \geq 0$, i.e. selling a very small number of shares potentially results in a non-zero cost. Furthermore, admissible trading times τ are assumed to be separated by $\Delta > 0$ units of time, i.e. $\tau_k = N \cdot \Delta$ for some integer N . Either of these assumptions is enough to rule out continuous trading policies. Hence we will set $S(t) = 0$ in (1), and focus on impulse type policies $\{\delta, \tau\}$. In the case where $\kappa = 0$, as $\Delta \rightarrow 0$ then we expect that these impulse policies should converge, in the proper sense, to the continuous trading policies discussed in the previous section.

The state space can be divided into two regions, a no-trade and a trade region. Let us define the trade operator \mathbf{T} as follows

$$\mathbf{T}[J(M, X, P, t)] = \sup_{\delta} \mathbf{E}[J(M_{t+\Delta}(\delta), X - \delta, P_{t+\Delta}(\delta), t + \Delta)]. \quad (19)$$

It is clear that trades occur when $J \leq \mathbf{T}[J]$.

The value function in this case will satisfy a differential inequality of the form

$$\max\left(J_t + \mathcal{L}[J], J - \mathbf{T}[J]\right) = 0. \quad (20)$$

The state space can be divided into the trade and no-trade regions as follows:

Trade:

$$J_t + \mathcal{L}[J] \leq 0, \quad J = \mathbf{T}[J];$$

No Trade:

$$J_t + \mathcal{L}[J] = 0, \quad J > \mathbf{T}[J].$$

The state space consists of two regions, rather than the three regions of the previous section, because of the nature of an impulse trading policy. In the case of continuous trading, being inside the trade region calls for a different action (a large sale) than being on the boundary of the trade region (an infinitesimal sale). However, in the case of discrete trading, being inside the trade region, or on its boundary with the no-trade region, requires a large (i.e. non-infinitesimal) sale.

Note that the revenue function R no longer needs to satisfy equation (16) when there are fixed costs of trading. In order to make specific statements about the optimal liquidation policy, we need to make an assumption about the time T utility function $u(\cdot)$ of the agent. We summarize our results for the costly trading case in the following theorem, the proof of which is also deferred to the Appendix.

Theorem 2 *Let $J(M, X, P, t) \in C^{2,2,2,1}$ be a solution to the differential inequality (20) with the boundary conditions (8-10). If J satisfies a set of regularity conditions, then $J(M, X, P, t)$ is the indirect utility function defined as in (7). Moreover, the optimal policy for X is to do nothing in the no-trade region and trade to the boundary or inside the no-trade region in the trade region.*

4 Risk Neutral Preferences with Continuous Trading

As a special case, let us assume that there are no transaction costs and that the investor is risk-neutral at time T . The revenue function in this case is given by (17). In this case, we are able to solve for the value function in closed form. The following theorem states the relevant results.

Theorem 3 For $J(M, X, P, T) = M + R(\delta, P)$, the value function is given by

$$J(M, X, P, t) = Me^{r(T-t)} + \frac{P}{\lambda}(1 - e^{-\lambda X})e^{\mu(T-t)}. \quad (21)$$

Furthermore, the condition in (15) has the same sign as

$$-(T - t)(\mu - r). \quad (22)$$

Proof. Simply check that the value function in (21) satisfies $J_t + \mathcal{L}[J] = 0$. (22) follows from simple algebra.

Q.E.D.

As can be seen from (22), as long as $t < T$ and the expected return on the stock is higher than the risk free rate, it is optimal for a risk-neutral agent to hold the stock until time T , and then to sell all X shares for a revenue of $R(X, P(T))$. If the expected return is less than the risk free rate, then it is optimal to sell all the shares right away.

5 Solutions for Power Utility Preferences

For the rest of the paper, we assume that the investor has a utility function with a constant relative risk aversion. Specifically, we assume that agents have utility over their time T money holdings given by

$$u(c) = \frac{1}{\gamma}c^\gamma. \quad (23)$$

We first study the case when there is no transaction costs for trading. In this case, the investor may trade both continuously and discretely. We then introduce transaction costs, which ensures that the investor will trade only discretely.

5.1 Power Utility with Continuous Trading

Given the power utility preferences, we conjecture that the value function has the following form

$$J(M, X, P, t) = M^\gamma g(\alpha, X, t), \quad (24)$$

where

$$\alpha(t) \equiv \log \left(\frac{P(t)}{M(t)} \right).$$

Recall that $R(\delta, P) = \frac{P}{\lambda}(1 - e^{-\lambda\delta})$. The boundary conditions in (8,9,10) become respectively

$$g(\alpha, X, T) = \frac{1}{\gamma} \left(1 + \frac{e^\alpha}{\lambda} (1 - e^{-\lambda X}) \right)^\gamma, \quad (25)$$

$$g(-\infty, X, t) = \frac{1}{\gamma} e^{r(T-t)\gamma}, \quad (26)$$

$$g(\alpha, 0, t) = \frac{1}{\gamma} e^{r(T-t)\gamma}. \quad (27)$$

It is then straightforward to check that

$$\begin{aligned} & \text{sgn}(J_M P - J_X - \lambda P J_P) \\ &= \text{sgn}(\gamma e^\alpha g - (\lambda + e^\alpha)g_\alpha - g_X). \end{aligned} \quad (28)$$

The Bellman-type equation in (11) becomes

$$\mathcal{L}[J] + J_t = g_t + r\gamma g + (\mu - r - \frac{1}{2}\sigma^2)g_\alpha + \frac{1}{2}\sigma^2 g_{\alpha\alpha}. \quad (29)$$

Hence we are able to reduce the problem from four state variables (M, X, P, t) , to three state variables (α, X, t) . This proves useful in solving the problem numerically.

5.2 Power Utility with Impulse Policies

Consider the case where a transaction cost is charged for each trade. Following standard practice in the literature (e.g. Morton and Pliska (1995) and Schroeder (1998)), we assume that $\kappa = k(M + XP)$ for some constant k . Hence a fixed fraction of an agent's wealth is charged every time a sell transaction is executed. Therefore

$$\lim_{\delta \downarrow 0} R(\delta, P) = -k(M + XP),$$

and $R(0, P) = 0$. Note that at the first glance, this form of transaction cost may appear unjustifiable. In reality, for liquidation in a short-horizon, the transaction cost κ is reasonably constant as the paper wealth $M + XP$ is reasonably constant. However, this choice of the function form makes it possible for us to reduce the dimension of state variables by one.

It is possible that at time T , the agent will be better off not to trade. We conjecture the same functional form as in (24), and note that the boundary conditions in (26,27) still apply. Hence the value function at time T is given by

$$g(\alpha, X, T) = \begin{cases} \frac{1}{\gamma} \left(1 + \frac{R(\delta, P)}{M}\right)^\gamma & \text{if } R(\delta, P) > 0 \\ \frac{1}{\gamma} & \text{if } R(X, P) \leq 0. \end{cases} \quad (30)$$

Note that the agent may decide to sell $\delta < X$ shares at time T . This is because for certain choices of revenue functions, we may have that $R(\delta, P) > R(X, P)$ for some $\delta < X$. Notice as well the requirement that $R(\delta, P)/M$ should be a function of α, δ, X, t only. The revenue function which we study is

$$\begin{aligned} R(\delta, P) &= \delta P e^{-\lambda\delta} - k(M + XP) \\ &= M \times \left(\delta e^{\alpha - \lambda\delta} - k(1 + X e^\alpha)\right). \end{aligned} \quad (31)$$

Notice that this implies that for a trade of δ shares, each share is sold at a price of $P e^{-\lambda\delta}$, where P is the pretrade price. It is this inability to sell each share at its marginal price (where the n th share is sold at a price of $P e^{-\lambda n}$) which provides an incentive to split up a large trade into smaller ones. Also that $R(\delta, P)$ achieves a maximum at $\delta = 1/\lambda$. Hence no per period sell amount will be

above this value. Notice also that our framework can handle any other form of revenue function which satisfies $R(\delta, P) = M \times f(\alpha, \delta, X, t)$.

To determine the no-trade region, we need to evaluate $\mathbf{T}[J]$. For the case of power utility function and a revenue function as specified above, we will have

$$\begin{aligned} \mathbf{T}[J(M, X, P, t)] &= \\ \sup_{\delta} M^{\gamma} \left(1 + \frac{R(\delta, P)}{M}\right)^{\gamma} g(\alpha', X - \delta, t), \end{aligned} \quad (32)$$

where

$$\begin{aligned} \alpha' &= \log\left(\frac{Pe^{-\lambda\delta}}{M + R(\delta, P)}\right) = \log\left(\frac{e^{\alpha}e^{-\lambda\delta}}{1 + R(\delta, P)/M}\right), \\ &= \alpha - \lambda\delta - \log(1 + R(\delta, P)/M). \end{aligned}$$

The sign of $J - \mathbf{T}[J]$ is therefore equal to

$$\text{sgn} \{g(\alpha, X, t) - (1 + R(\delta, P)/M)^{\gamma} g(\alpha', X - \delta, t)\}.$$

Note that since trade occurs only if $R(\delta, P) > 0$, we have $\alpha' < \alpha$. Also we know that $\partial g/\partial\alpha \geq 0$ and $\partial g/\partial X \geq 0$. Therefore we can establish that $g(\alpha, X, t) \geq g(\alpha', X - \delta, t)$.

6 Comparisons with the Merton Problem

In this section we compare our model with that of Merton (1971) with no price impact. We first show that if we allow the investor to buy and sell shares as he wishes, then the optimal solution method introduced in the previous sections reduces to the standard Merton's solution. To build intuition about our results, we then solve the Merton problem with the restriction that the investor only can sell his holding. In both cases, we assume that the investor's trades do not move the prices.

6.1 Equivalence to the Standard Merton's Solution

In this section, we establish that if no trading frictions exist, and if both sales and purchases are allowed, then our framework, which allows the investor to control the process X based on information of X and M , reduces to the standard Merton's portfolio choice problem.

The dynamics of the state variables in the no market impact case are as follows

$$\begin{aligned} X(t) &= X(0) + \int_0^t l(s)ds + \int_0^t h(s)dw(s), \\ M(t) &= M(0) + \int_0^t rM(s)ds + \int_0^t (P(s) + dP(s))dX(s), \\ P(t) &= P(0) + \int_0^t \mu P(s)ds + \int_0^t \sigma P(s)dw(s). \end{aligned}$$

In order to be consistent with the Merton's setup, we allow the investor to trade shares both deterministically (l) and stochastically ($h dw$). However, in contrast to the Merton's problem

where the dynamics of X and M are reduced to the single wealth dynamics, the investor is allowed to simultaneously choose l and h to control X and M . As usual, we will only consider feedback controls with l and h being functions of state variables X , M , P and time t .

To prove that the optimal solution of our problem reduces to that of the Merton problem, we take the Merton solution as given and define the following:

$$\widehat{X}(t) = \widehat{\pi} \left(\widehat{M}(t) + \widehat{X}(t)P(t), t \right) \frac{1}{P(t)}$$

where $\widehat{M}(t)$ and $\widehat{X}(t)$ denote the money market and share holdings corresponding to the Merton's solution,⁶ $\widehat{\pi}$ denotes the optimal proportion invested in the risky stock (as a function of wealth and time t) obtained from the Merton's solution. We want to show that the \widehat{l} and \widehat{h} corresponding to $d\widehat{X}$ defines the optimal solution to our problem. Our proof below is largely heuristic and can be made rigorous by imposing the regularity conditions wherever necessary. For a rigorous treatment and technical details, we refer readers to Hindy, Huang and Zhu (1997) in which a similar problem was solved.

Suppose that h^* and l^* define the pair of optimal policies. Let us consider a small perturbation along the optimal policies l^* and h^* : for a small $\Delta t > 0$,

$$l_1(s) = \begin{cases} l(\cdot, s), & t \leq s \leq t + \Delta t \\ l^*(X(s), M(t), P(s), s), & t + \Delta t < s \leq T \end{cases}$$

$$h_1(s) = \begin{cases} h(\cdot, s), & t \leq s \leq t + \Delta t \\ h^*(X(s), M(t), P(s), s), & t + \Delta t < s \leq T \end{cases}$$

where $l(\cdot, s)$ and $h(\cdot, s)$ are some arbitrary policies, i.e., the investor switches to the optimal policies from $t + \Delta t$. Let J_1 denote the indirect utility function given policies l_1 and h_1 and J^* the indirect utility function given policies l^* and h^* . Then, by the rule of iterative conditional expectation, we have

$$J_1(M(t), X(t), P(t)) = \mathbf{E}_t \left[J^*(M(t + \Delta t), X(t + \Delta t), P(t + \Delta t), t + \Delta t) \right] \quad (33)$$

Next, applying Ito's lemma to J^* (which is a function of M, X, P and t), we find that

$$J^*(M(t + \Delta t), X(t + \Delta t), P(t + \Delta t), t + \Delta t) = J^*(M(t), X(t), P(t), t) + \int_t^{t+\Delta t} \psi(s) ds$$

$$+ \int_t^{t+\Delta t} J_M^*(s) P(s) h(s) dw(s) + \int_t^{t+\Delta t} J_X^*(s) h(s) dw(s) + \int_t^{t+\Delta t} J_P^*(s) \sigma P(s) dw(s)$$

where

$$\psi = J_t^* + \mu P J_P^* + \frac{1}{2} \sigma^2 P^2 J_{PP}^* + r M J_M^* + l(P J_M^* - J_X^*)$$

$$+ h P \sigma (J_{XP}^* - J_M^* - P J_{MP}^*) + h^2 \left(\frac{1}{2} J_{XX}^* - P J_{MX}^* + \frac{1}{2} P^2 J_{MM}^* \right)$$

Plugging the above expression for $J^*(t + \Delta t)$ into (33) and noting that $J_1 \leq J^*$, we conclude that⁷

$$\mathbf{E}_t \left[\int_t^{t+\Delta t} \psi(s) ds \right] \leq 0 \quad (34)$$

⁶To make it rigorous, we shall define $\widehat{M}(t)$ and $\widehat{X}(t)$ path by path.

⁷Regularity conditions are required to make sure that all Ito integrals have zero expectation.

We note that since there is no price impact, the investor can always instantaneously convert shares into money or vice versa. As a result, we can argue that the indirect utility function along the optimal path must be a function of wealth and time t . Let us define wealth as $W = M + XP$ and

$$J^*(M, X, P, t) = V(M + XP, t)$$

It is easy to verify that in this case

$$\begin{aligned} PJ_M^* - J_X^* &= 0, \\ J_{XP}^* - J_M^* - PJ_{MP}^* &= 0 \\ \frac{1}{2}J_{XX}^* - PJ_{MX}^* + \frac{1}{2}P^2J_{MM}^* &= 0. \end{aligned}$$

This simplifies ψ to as follows,

$$\psi = V_t + \mu PXV_W + \frac{1}{2}\sigma^2P^2X^2V_{WW} + rMV_W$$

For $t \leq s < t + \Delta t$, we now set

$$\begin{aligned} l(\cdot, s) &= \widehat{l}(s) \\ h(\cdot, s) &= \widehat{h}(s) \end{aligned}$$

Notice that by definition, $\widehat{\pi}(s) = \widehat{X}(s)P(s)/\widehat{W}(s)$ where $\widehat{W}(s) = \widehat{M}(s) + \widehat{X}(s)P(s)$, we have $P\widehat{X}(s) = \widehat{\pi}(s)\widehat{W}(s)$ and $\widehat{M}(s) = (1 - \widehat{\pi}(s))\widehat{W}(s)$. Therefore we can write ψ as

$$\psi(s) = V_t(\widehat{W}(s), s) + (r + (\mu - r)\widehat{\pi})\widehat{W}(s)V_W(\widehat{W}(s), s) + \frac{1}{2}\sigma^2\widehat{\pi}^2(s)\widehat{W}(s)^2V_{WW}(\widehat{W}(s), s)$$

which is zero since $\widehat{\pi}(s)$ is optimal for the Merton's problem and the Bellman equation for the Merton problem holds. We conclude from (34) that

$$J_1(M(t), X(t), P(t), t) \leq J^*(M(t), X(t), P(t), t)$$

with equality holds when $l(\cdot, t) = \widehat{l}(t)$ and $h(\cdot, t) = \widehat{h}(t)$. Hence, \widehat{h} and \widehat{l} are indeed optimal. We conclude that our solution coincides with Merton's solution.

6.2 Optimal Liquidation with No Price Impact

As a comparison, we now analyze the optimal liquidation problem when there is no price impact, i.e., $\lambda = 0$. This case differs from the standard portfolio problem because only sales, and not purchases, are allowed. Because there is no price impact, knowing α and X separately is unnecessary for the formulation of this problem. Let us define a new state variable $\xi \equiv \log(XP/M)$. We can then write the value function as $J(M, X, P, t) = M^\gamma h(\xi, t)$. From (24) we have $g(\alpha, X, t) = h(\alpha + \log X, t)$, and then boundary conditions (25, 26, 27) imply that $h(\xi, T) = 1/\gamma(1 + \exp(\xi))^\gamma$ and $h(-\infty, t) = 1/\gamma \exp(r(T - t)\gamma)$. From (28) we have that $\text{sgn}[J_M P - J_X - \lambda P J_P] = \text{sgn}[\gamma h - h_\xi(\exp(-\xi) + 1)]$. And finally from (29) we have

$$\mathcal{L}[J] + J_t = h_t + r\gamma h + (\mu - r - \frac{1}{2}\sigma^2)h_\xi + \frac{1}{2}\sigma^2 h_{\xi\xi}.$$

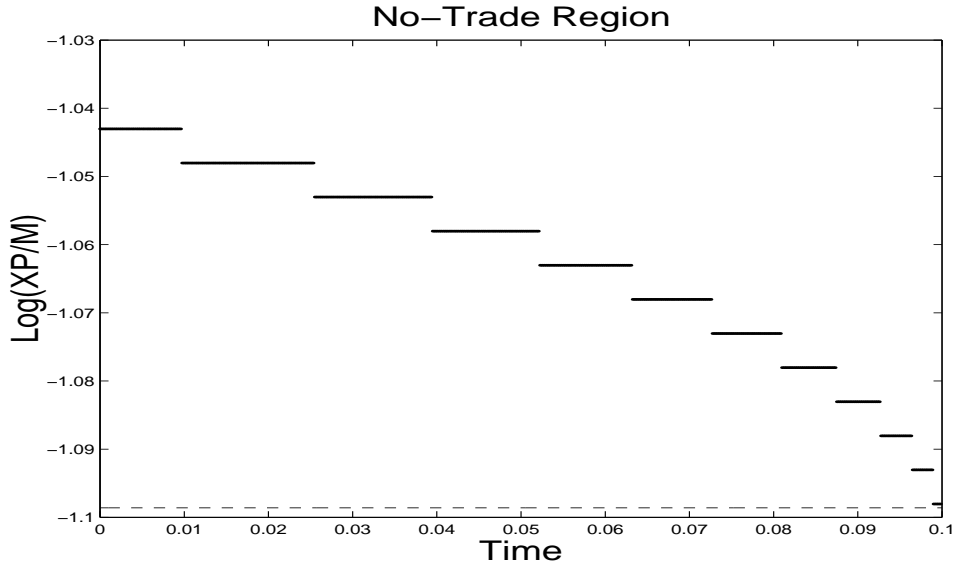


Figure 1: The dashed line is the optimal solution for the standard Merton problem. The solid line is the no-trade region. The parameters values are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0$, $r = 0.05$, $T = 0.1$, $X(0) = 10$, $A(0) = 2$, $\gamma = -3$, $k = 0$.

Unfortunately no closed form solution for this problem exists. We therefore solve it numerically. Figure 1 shows the no-trade region for the standard Merton problem with no price impact. The dashed line shows the optimal level of ξ in the standard Merton problem where buys and sells are allowed (i.e. $\hat{\pi} = (\mu - r)/((1 - \gamma)\sigma^2)$ and $\exp(\hat{\xi}) = \hat{\pi}/(1 - \hat{\pi})$). The solid line shows the no-trade region as a function of ξ and t (the discontinuities are due to discreteness of the numerical solution). Hence the agent sells just enough shares to maintain ξ inside the no-trade region at the time when ξ hits the boundary.

The boundary for the no-trade region is above $\hat{\xi}$, indicating that the agent's optimal holdings of the risky asset is to be controlled above $\hat{\xi}$ – the optimal level corresponding to the Merton's problem. The basic intuition for this result is as follows. The agent would like to ensure that his holding ξ remains close to the unconstrained optimal level $\hat{\xi}$ on average. This can be best achieved by controlling ξ to be slightly above $\hat{\xi}$. When $\xi > \hat{\xi}$, the agent is willing to tolerate this situation because if the stock price falls, ξ will fall closer to $\hat{\xi}$, whereas if the stock price rises, the agent can decrease ξ by selling more shares. On the other hand, if $\xi < \hat{\xi}$, then a drop in stock price may move ξ to be too far away from $\hat{\xi}$ since the agent is not allowed to increase ξ by buying more shares. The impact of the no-purchasing restriction is more substantial the further away it is from the final date. That is why the boundary for the no-trade region is decreasing in time to maturity. However, as t approaches T the probability that a ξ which is far above $\hat{\xi}$ will ever fall below $\hat{\xi}$ is very small. Therefore the amount by which ξ may exceed $\hat{\xi}$ decreases to zero as t approaches the consumption date T . Finally, at time T the agent sells all of his share holdings for a revenue of $X(T) \cdot P(T)$.

7 Numerical Results: The General Case

In the general case of the model which we solve numerically, agents are able to trade every Δ units of time, where Δ corresponds to the discretization in time. Additionally, every trade may incur a transaction cost proportional to the agent's total wealth at the time of trade. In this section, we analyze the optimal sale of a large block of shares under these circumstances.

Recall from (31) that in the case of discrete trading each share in an order of δ shares is sold at a price of $P \exp(-\lambda\delta)$. Because of the convexity of the price impact function, it is better to sell 1 share followed by 1 share than to sell 2 shares right away (i.e. $e^{-\lambda} + e^{-2\lambda} > 2e^{-2\lambda}$). However, since trading is allowed to occur only every Δ units of time, splitting up an order potentially carries the cost of being away from the optimal portfolio allocation while economic uncertainty is being realized.

For example, if the stock's returns are very volatile, then failing to sell shares in the current trading period may lead to very large losses in the next trading period. On the other hand, it may be advantageous to delay sales of shares if the expected return on the stock is higher than the return earned by the money account. Clearly, the tradeoff of whether to sell early or late depends on the stock's return process, the risk free rate, and the agent's risk preferences. Furthermore, if a transaction cost is charged every time that an order takes place, then the incentive to engage in multiple transactions diminishes.

7.1 Discussion of Parameter Values

We will consider the case of an agent trying to sell 10 units of a security, with a price impact coefficient of $\lambda = 0.01$. Furthermore, we normalize the initial price to 1, i.e. $P(0) = 1$, and the initial money account to $M(0) = \exp(-2)$. Hence the agent's initial wealth of $M(0) + X(0)P(0) = 10.135$ consists mostly of the stock position. Note that immediate liquidation of this 10 share block would result in a post trade price of $\exp(-0.01 \times 10) = 0.9048$, or a price impact of approximately 10%. This choice of parameters is simply a normalization made for expositional convenience. If we interpret a unit of security as representing 100,000 shares of some stock with a price of \$10 and re-scale λ to 10^{-6} , then the order which we are considering consists of selling \$1,000,000 worth of a security in an environment where the immediate liquidation of all these shares would move the price to \$9. Also, given this parameterization, the agent's initial holdings of money is \$13,534.

Recall that when $k > 0$ a fixed cost proportional to the agent's current wealth (i.e. $\kappa = k(M + X P)$) is charged every time a trade is made. Note firstly that the quantity $M + X P$ is not expected to change very much over the short time horizon (e.g. $T = 0.1$) which we are considering. Hence the actual cost paid κ is expected to remain fairly constant. For those simulations where a transaction cost is charged, we assume a cost of 10 basis points of the portfolio value, or $k = 0.001$. Hence the per trade cost would be approximately $0.001 \times \$1,000,000 = \$1,000$ in this example. This is the cost which must be paid every time an order is executed. Components of this fixed cost include ticket charges, minimum transaction fees, and the search costs of working an order. It is likely that such fixed transactions costs are more prevalent in illiquid or emerging markets than they are on extremely liquid exchanges such as the NYSE.

Note as well, that the fixed per trade cost does not include any per share costs associated with transactions. Per share costs are analogous to the price impact associated with trading, and hence would not markedly affect optimal trading behavior. However, per transaction costs introduce an incentive to economize on the number of transactions which should be performed. As we show this incentive can have a pronounced affect on optimal policies.

7.2 Optimal Liquidation with Continuous Trading

Before we turn to the analysis of the general case, let us consider optimal liquidation in the case where transaction costs are 0 while continuous trading is allowed. This case extends the model of Section 6.2 by introducing price impact. We find a behavior similar to the no-price impact case, with an adjustment to the no-trade region for the fact that sales cause the price to fall.

Recall that the three relevant state variables are now $\alpha = \log(P/M)$, X , t . Figure 2 shows the no-trade regions for two cases of the model, one with a higher price impact term λ than the other. From Figure 2, we see that the no-trade region is decreasing as a function of X . In other words, the investor is more willing to unload his shares when he has a larger position. We also observe that the no-trade region at $t = 0$ for the low price impact case is lower than for the high price impact case. Note that for a given share holding X , a higher no-trade region implies a willingness to keep a higher value of the ratio between the investment in the risky asset (XP) and investment in the riskless asset (M). This result shows that investors facing larger market impact generally try to unload their share holdings at a slower rate so as to keep the ratio XP/M at a higher level. This is again related to the restriction that investors can only sell shares. With larger market impact, an aggressive liquidation of shares may cause price to drop too much, which may push the ratio too much away from the Merton's line. This is beneficial because by holding on to shares for a longer time period, the expected price at which shares will be sold is higher, and this is needed to counteract the effect of a larger price impact from sales.

Since sales take place at marginal prices, no incentive exists for splitting up the block into multiple smaller transactions. The no-trade region falls slightly, and maintains its general shape, as t approaches T (as in Figure 1) for the same reasons as in the $\lambda = 0$ case. At time T all remaining shares are liquidated for a revenue of $\frac{P(T)}{\lambda}(1 - e^{-\lambda X(T)})$. Note that for the parameter values in Figure 2, a sale of 10 shares with a price impact of $\lambda = 0.01$ would result in revenue of $9.5163 \times P$, and a post-trade price of $e^{-\lambda X} P = 0.9048 P$.

We expect that in the discrete trading case, when the transactions costs for trading are 0, that as $\Delta \rightarrow 0$ the optimal trading policy should approach the one in the continuous trading case: agents will optimally split up the block of shares into many small transactions, so as to approximate as closely as possible the liquidation strategy in the continuous trading case. In particular, even with no transactions costs, we expect that a no-trade region will arise in the impulse trading case as Δ becomes sufficiently small. As Figure 3 shows, the no-trade region in the $\Delta > 0$ case does indeed approach the no-trade region in the case of continuous trading. The circles in the figure represent the no-trade region in the latter case, the x's represent the no-trade region in the case

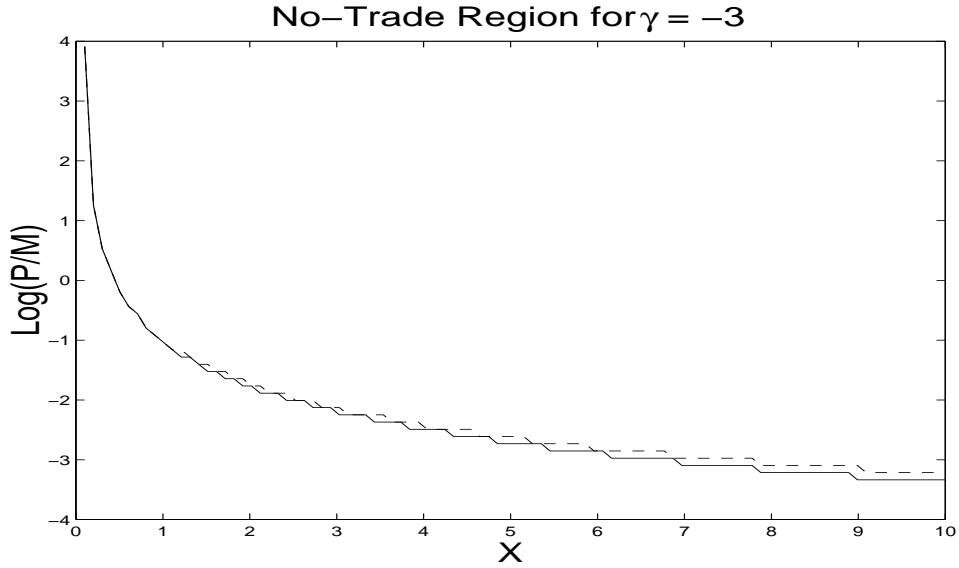


Figure 2: The parameters values for the solid line are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$, $\gamma = -3$, $k = 0$. The dashed line has $\lambda = 0.04$.

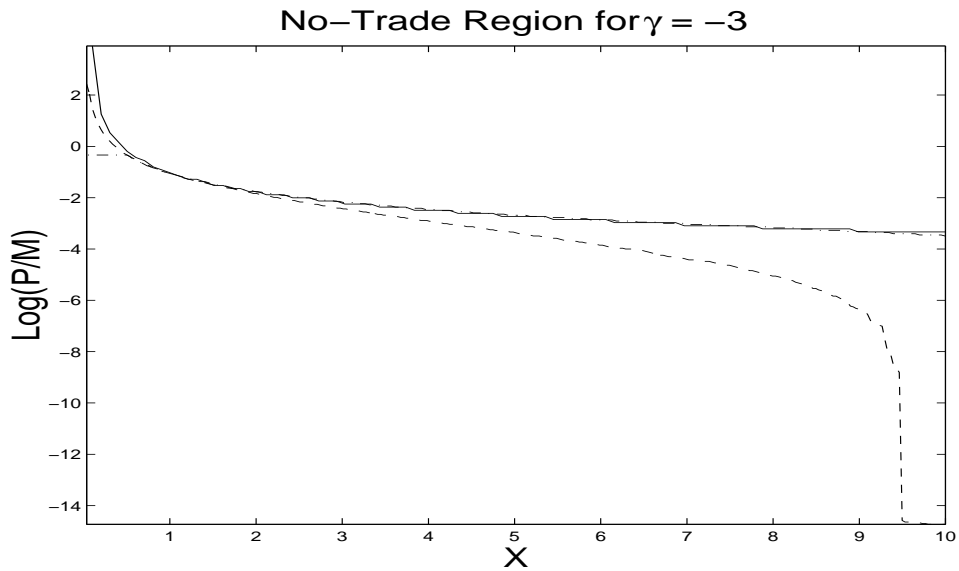


Figure 3: The time 0 no-trade regions under different policies. The solid line represents the no-trade region in the case of continuous trading. The dashed-dotted line (almost indistinguishable from the solid line) represents the $\Delta > 0$ case with $N_T = 75$, and the dashed line corresponds to $N_T = 20$, where N_T is the number of trades allowed between times 0 and T . The parameters values are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$, $\gamma = -3$, $k = 0$.

where $\Delta = 0.1/20$, and the dots correspond to the case where $\Delta = 0.1/75$.⁸

7.3 Revenue Distributions

In our setup agents care about utility over their money account at some future time T . Hence we can gain insight into the properties of optimal liquidation policies by studying the distribution of $M(T)$ (i.e. the time T money account), as well as the quantity $\mathbf{E}[u(M(T))]$. While the latter quantity is simply equal to the agent’s value function, the former distribution is difficult to compute directly. Hence we will bootstrap the distribution of $M(T)$ under various liquidation policies by simulating the model.

In order to provide a baseline for understanding our results, we will assume that agents can alternatively use the “naive” policy of evenly splitting up their trades over available trading rounds. Under this policy, an agent attempting to sell X shares over N trading rounds would choose to sell $\delta = X/N$ shares in each round. In fact, Lo and Bertsimas (1998) show that under some circumstances this “naive” policy is, in fact, optimal.⁹ Hence it provides a natural benchmark for the evaluation of the optimal policies under our approach.

In the subsequent analysis, instead of focusing on the variable $M(T)$ it is more convenient to see what return on his initial wealth the agent can achieve under different liquidation schemes. Hence we define $R(T)$ as the simple net return under a given liquidation strategy, or

$$R(T) \equiv \frac{M(T)}{M(0^-) + X(0^-)P(0^-)} - 1.$$

This measures the return on the agent’s initial wealth. Because of the price impact of selling, we expect that $R(T)$ will be negative over short time horizons. Another interesting characteristic of a trading strategy is the share weighted price at which the order was sold. We define $\Pi(T)$ as

$$\Pi(T) \equiv \frac{\sum_{s=0}^T \delta(s) \cdot P(s)}{P(0^-) \cdot \sum_{s=0}^T \delta(s)}.$$

This is the revenue from all sales, normalized by the total number of shares sold times the initial price. Previous work in the area of optimal block liquidation has focused mainly on properties of $\Pi(T)$. Lo and Bertsimas (1998) study the control problem of an agent with the goal of maximizing $\mathbf{E}[\Pi(T)]$. Huberman and Stanzl (2001) study the problem of an agent who wants to maximize $\mathbf{E}[\Pi(T)] - \mathbf{Var}(\Pi(T))$. Subramanian and Jarrow (2001) analyze the related problem of an agent who maximizes $\mathbf{E}[\sum_t u(\delta(t) \cdot P(t))]$, where $u(\cdot)$ is some utility function. In our case, of course, $\Pi(T)$ is simply a descriptive statistic, rather than an objective function.

We will consider the following four cases in our analysis: (1) the case of a risk averse individual and no transactions costs ($\gamma = -3$ and $k = 0$), (2) the case of a risk neutral individual with no transactions costs ($\gamma = 1$ and $k = 0$), (3) the case of a risk averse individual with transactions costs

⁸The horizontal line in the upper left hand part of the figure is a consequence of the numerical imprecision of the solution, and is not a feature of the no-trade region.

⁹Specifically, Lo and Bertsimas assume that risk-neutral agents face a linear price impact function, and a stock price with a 0 expected return.

($\gamma = -3$ and $k = 0.001$), and (4) the case of a risk neutral individual with transactions costs ($\gamma = 1$ and $k = 0.001$).

Tables 1, 3, 5, and 7 show summary statistics of simulations of the model for cases (1–4) respectively. Variables marked with a \cdot^* refer to variables under the optimal policy, and unmarked quantities refer to the naive, linear strategy. The tables report the means, standard deviations, and percentiles of $R(T)$ and $\Pi(T)$ across the simulation runs. The bottom part of the tables reports the mean and standard error of $u(M(T))$, and $J(0^-)$ is the agent’s initial value function realization.

As can be seen from Table 1, for a risk averse agent with no transactions costs, the average return from the optimal policy (measured in terms of $R(T)$ or $\Pi(T)$) is actually lower than the return from the naive strategy. However, the average utility (i.e. the mean of $u(M(T))$) for the optimal policy is higher than for the naive policy. The reason for this lies in the fact that the optimal policy achieves a much less variable distribution of both $R(T)$ and $\Pi(T)$ than does the naive strategy. This can be seen both from the reported standard deviations, and from the percentiles of the bootstrapped distributions for the two variables. Hence in order to achieve an ex-ante less variable strategy, a risk averse agent is willing to give up something in the expected return from the liquidation.

A risk neutral agent, on the other hand, should be able to achieve a higher expected return than the naive strategy, since maximizing the expected $M(T)$ is exactly equivalent to maximizing $R(T)$. As Table 3 shows a risk neutral agent, facing no transactions costs, is able to achieve a higher return and higher share weighted price than under the naive strategy. However, this higher return comes at a price (unimportant for a risk neutral agent) of having a higher variance of realized revenue. From the standard deviations and from the percentiles of the distribution, we can see that the optimal risk neutral strategy results in much more variable realizations of $R(T)$ and $\Pi(T)$ than does the naive, linear strategy.

As these two tables suggest, in the case of no transactions costs a trade off exists between liquidation strategies which do well in terms of expected returns and those which do well in terms of minimizing the variability of realized revenues. Which of these considerations dominates is a function of agents’ preferences for risk. The reason why these tradeoffs exist will be discussed further in the next section.

As may have been expected, in the presence of transactions costs, the naive strategy does extremely poorly relative to the optimal strategies. The reason for this is that the naive strategy overtrades, and hence results in overly high transactions costs. The optimal strategies, by prescribing no trade during some of the trading times, economize on transactions costs. Table 5 shows that for the optimal risk averse strategy, the average returns $R(T)$ are higher than those under the naive strategy. The revenue variability is also lower under the optimal strategy. Table 7 compares the optimal risk neutral strategy to the naive strategy. The risk neutral strategy achieves higher average returns $R(T)$ than both the optimal risk averse and the naive strategies. However, the variability of revenues is also highest under the optimal risk neutral strategy.

Notice that with transactions costs, while the optimal strategies achieve higher average returns, they also achieve lower share weighted prices (i.e. $\Pi^*(T) < \Pi(T)$). This suggests that minimizing transactions costs comes at the expense of not minimizing the price impact of trading. Since

$\Pi(T)$ does not take into account transactions costs charged against the money account, it is an inappropriate measure of the goodness of a policy when costs are involved.

7.4 Properties of the Optimal Policy

In order to better understand the results in the previous section, let us consider the structure of an optimal trading policy. Figure 4 shows the no-trade regions for a risk averse and risk neutral agent as a function of time left until T . In the figures the x-axis represents the number of shares held by the agent, and the y-axis represents α . The no-trade region lies to the lower left hand side of the boundary plotted in the graphs. As can be seen, the no-trade regions move towards the left as t approached T . For example, the dashed line in both graphs shows the boundary of the no-trade region at time $t = 0.25 T$. If the values of the agent's state variables X, α lie in the region below and to the left of the boundary then the agent would not trade at time t . On the other hand, if X, α lie above or to the right of the boundary the agent would trade. As t approaches T the size of the no-trade region decreases. This implies that the agent slowly sells off shares of the risky asset as time T approaches.

Up to numerical error, the rightmost part of the no-trade region for a risk neutral agent is vertical. This implies that the risk neutral agent does not take into account α when deciding whether or not to trade. Since $X \cdot e^\alpha = XP/M$ is a measure of how risky an agent's portfolio is, the fact that the risk neutral agent does not take this ratio into account is not surprising – after all, in the absence of trading frictions the risk neutral agent's optimal portfolio is to invest an infinite amount of money into whichever asset has the highest expected return.

On the other hand, holding X and t fixed, the risk averse agent will choose to sell shares for a sufficiently high value of α . For a high enough α , the risk averse agent will hold too much of his wealth in the risky asset, and will therefore choose to sell off some of it. Whereas the risk neutral agent's policy depends only on t and X , the risk averse agent also takes into account how much risk exposure he has in deciding the speed at which to liquidate his portfolio.

Figure 5 shows the optimal trade amounts for a risk averse agent for values of X, α which are in the trade region. The optimal trade amount for a risk averse agent is increasing in both X and α . The optimal trade amount is increasing in X because with more shares held, more shares need to be sold per trading opportunity in order to achieve the minimum possible price impact. The trade amount is increasing in α because α measures the amount of risk exposure which the agent faces at a given X . The higher the risk exposure, the more of the risky asset a risk averse agent would like to sell. Though it is not shown here, the analogous surface for a risk neutral agent is increasing in X , but does not depend on α .

Another perspective on the optimal liquidation policy prescribed by our approach is to examine the time series properties of the optimal per period trade amounts. The solid lines in Figure 6 show the average δ in each trading period across the simulations discussed in the previous section. The dashed lines represent two standard deviation bounds around the mean trade amount in each period.¹⁰ For the risk averse agent facing no transactions costs, we see that the optimal liquidation

¹⁰It should be noted that the standard deviation bounds appear to be due mostly to imprecision of the numerical

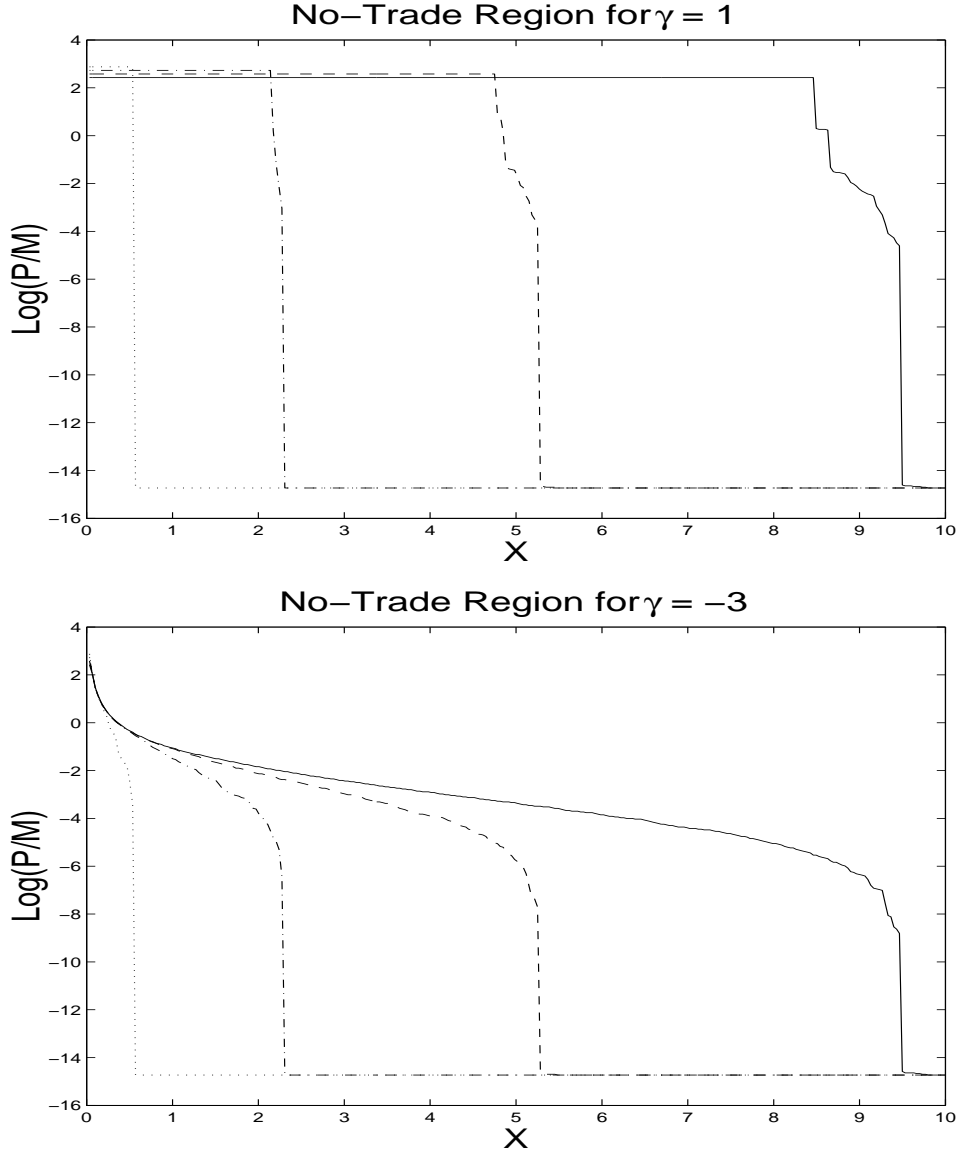


Figure 4: The no-trade region boundaries plotted for different time periods. The regions themselves are to the left of the boundaries. The solid, dashed, dashed-dotted, dotted lines correspond respectively to 20, 15, 10, 5 trading periods remaining before time T . The parameters are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$, $\kappa = 0$.

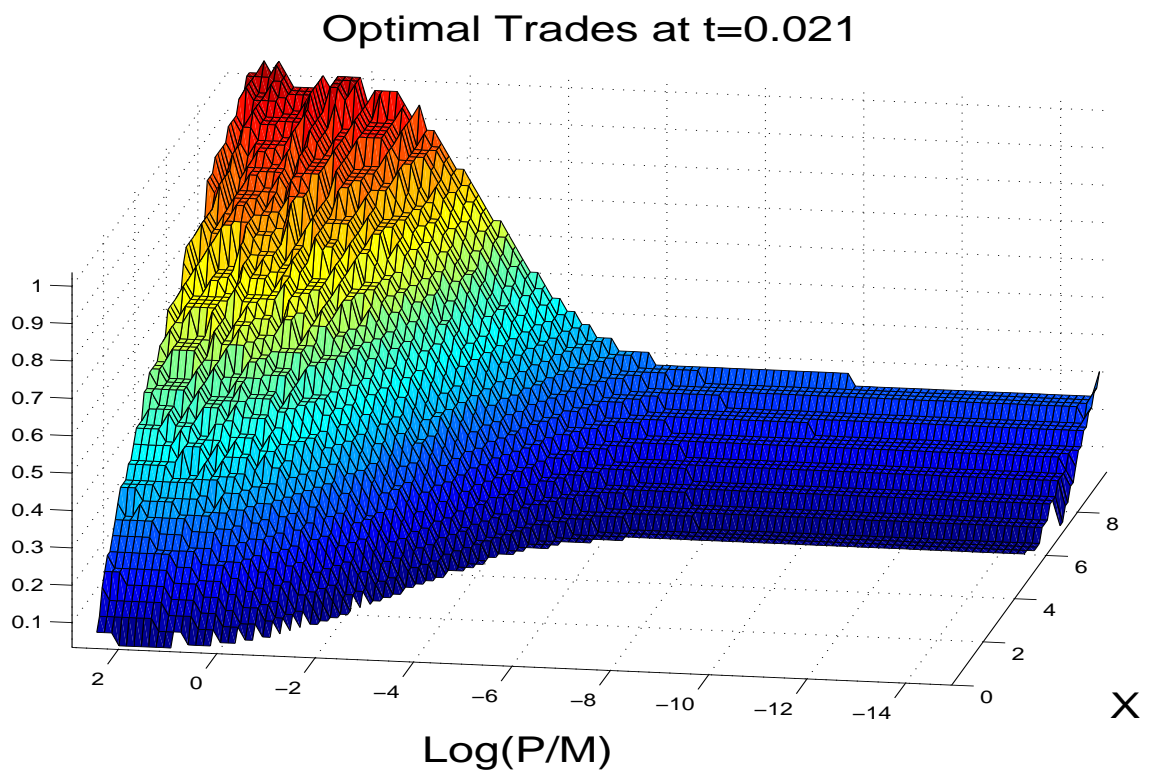


Figure 5: Optimal trade sizes. The parameter values are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$, $\gamma = -3$, $k = 0$.

policy is u-shaped. Hoping to decrease the variability of his time T money holdings, the risk averse agent engages in heavy initial selling. At some point the agent reaches a desirable portfolio from a risk-return point of view, and begins to sell more heavily only when few trading periods remain before time T .

The risk neutral agent displays a very different trading behavior. He sells very little at the beginning, preferring to stay exposed to the risky asset for as long as possible since it offers a higher expected return than the risk free rate ($\mu > r$). Compared to the policy of the risk averse agent, this trading behavior leads to a higher expected value of $M(T)$, but also to a higher variance of $M(T)$. Since the stock has a higher expected return than the money account, by holding on to the stock for longer, the risk neutral agent is able to achieve a higher expected return from liquidating the portfolio. However, this strategy leaves the agent more exposed to price fluctuations. Hence a tradeoff exists between the expected return and the volatility of a liquidation strategy. Even though the time horizon of relevance to this problem is quite short, perhaps a month or less, the gain in expected returns from the risk neutral strategy can still be significant. Given the parameterization used in Table 3 it is around 10 basis points per 0.1 year, or about 1 percent per year.

In the presence of transactions costs, the general shape of each agent's policy remains the same. However, agents transact on fewer occasions, choosing instead to sell more shares in those times when they do trade. While this strategy incurs a higher amount of price impact, it minimizes the amount of transactions costs paid by the agents. A risk averse agent sells many shares initially, and then does not trade until the very last periods. The risk neutral agent does not sell at all at the beginning, and only begins to sell as time T draws near.

Tables 2, 4, 6, and 8 show the price, sale amount, shares held after trading, and money account balance for each trading round in a representative simulation run. The tables correspond, respectively, to cases (1–4) discussed in the previous section. The variables marked with a \cdot^* correspond to the optimal strategy, and the unmarked variables correspond to the naive, linear strategy. Under the latter, it can be seen from the tables that δ is the same in every trading period. The tables tell a similar story to Figure 6. Note that similar to the result in (3), the optimal and naive strategies lead to an identical final price.

7.5 Liquidity Discount

Consider an optimizing agent who is given a block of shares of the risky security. How much money would that agent demand in order to part with that block of shares? This amount should reflect the expected revenue to the agent from optimally liquidating that block, as well as the disutility associated with the uncertainty of the liquidation process. This money amount is a much more accurate measure of the “worth” of those securities than their current market price multiplied by the number of shares in the block. This *certainty equivalent* amount, C , solves the following equation

$$J(M, X, P, t) = J(M + C(X), 0, P, t),$$

solution (see also Figure 5). Over the short time horizon ($T = 0.1$) which we are considering, not much variation exists in the optimal liquidation path for our chosen set of starting values for X, P, M .




Figure 6: Shares traded per time period across all simulation runs. The solid line shows the mean per period trade amount across all simulations. The dashed lines show two standard deviation bounds around the per period mean trade amount. The parameters are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$.




Figure 7: Per share prices as a function of X . The parameter values for the solid line are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$, $\gamma = -3$, $k = 0$. The dashed line has $\sigma = 0.2$ and $\lambda = 0.011$, and the dashed-dotted line has $\sigma = 0.3$, $\lambda = 0.01$ and $\gamma = 1$.

where $C(X)$ denotes explicitly the dependence on the number of shares in the block.¹¹ We can then think of $C(X)/X$ as that agent's shadow price for X shares of the stock. If we let $M = 1$ and use the functional form for J from (24) then we can write this as

$$(1 + C(X))^\gamma \frac{e^{r(T-t)\gamma}}{\gamma} = g(\log P, X, t).$$

We can then solve for $C(X)$ to find that

$$C(X) = \left(e^{-r(T-t)\gamma} \gamma g(\log P, X, t) \right)^{1/\gamma} - 1. \quad (35)$$

One measure of a security's liquidity is the per share discount (or premium) relative to the market price at which an optimizing agent would be willing to part with a large block of shares. The agent would be willing to accept a discount relative to the current share price because of the riskiness involved with liquidating the position, as well as because of the price impact which such liquidation would incur. Let us define this price discount as

$$l(X) \equiv \frac{C(X)}{X \times P}. \quad (36)$$

Figure 7 plots $l(X)$ for different parameter values of the model.

Given existing market conditions, to a seller of a large block of shares a more liquid market implies that $l(X)$ is higher – that is, the seller demands a per share amount closer to the actual market price of the security. Hence $l(X)$ is one measure of the “effective liquidity” of the market. It turns out that $l(X)$ is significantly affected by parameters of the model other than λ , which is the traditional measure of market liquidity used in both academic and industry circles. The solid line in Figure 7 shows $l(X)$ for $\lambda = 0.01$ and $\sigma = 0.3$, and the dashed line shows $l(X)$ for $\lambda = 0.011$ and $\sigma = 0.2$. As can be seen, the effective liquidity of the market is actually higher for the stock with the larger price impact of $\lambda = 0.011$ for sales of 4.5 shares or less. For sales higher than this, the effective liquidity is higher for the stock with the smaller price impact of $\lambda = 0.01$.

There are two effects in the way that such a trade is priced by the risk averse agent. First, those securities which are more desirable from a risk-return point of view fetch higher prices. Hence the security with the lower σ is preferable when sufficiently few shares have to be liquidated and hence the price impact is relatively small. Indeed, for a sufficiently small number of shares to sell, $l(X) > 1$ indicating that the agent would demand more than the current market price to part with his shares – this is a consequence of the fact that the agent is allowed only to sell, and not to buy, shares. Once enough shares of a security need to be liquidated, the price impact effect dominates and the risk-return characteristics of the security matter less in the way the order is priced.

Notice that the risk preferences of the agent also have an important effect. For the case of a risk-neutral agent, the effective liquidity is the highest possible (for a given μ, r, λ) because the

¹¹Another interpretation would be that the trading desk is willing to pay $C(X)$ dollars to a client wanting to sell it X shares (i.e. $J(M, 0, P, t) = K(M - C(X), X, P, t)$). The desk would then optimally liquidate the X shares after having paid out $C(X)$ to its client. The client would presumably accept $C(X)$ dollars so as to avoid the uncertainty of going into the market and selling those shares by himself over some time period. Hence $C(X)/X$ would be the per share price which the client could expect to get for the shares immediately by going through the trading desk.

volatility of the stock's price does not matter to the agent. The dashed-dotted line in Figure 7 shows the $l(X)$ for a risk-neutral agent and a price impact term of $\lambda = 0.01$.

An interesting comparison can be made between $l(X)$ and $\Pi(T)$. From Figure 7, we see that for the risk averse agent $l(X = 10)$ is approximately 0.945, whereas from Table 1 we see that the mean value of $\Pi^*(T)$ is 0.954. Hence a risk averse agent accepts a fairly substantial discount to the expected share weighted liquidation price (keep in mind that $\Pi(t)$ ignores the fact that sales proceeds are reinvested at the riskless rate until time T , and hence underestimates the true expected revenue per share). For the risk neutral agent, $l(X = 10) \approx 0.954$, and from Table 3, $\Pi^*(T) \approx 0.958$. The discount demanded by the risk neutral agent reflects the present value of anticipated future sales proceeds, and hence reflects the anticipated price impact of the block liquidation. The discount demanded by the risk averse agent includes an additional risk discount.

7.6 Simple Extensions

Our model can be generalized in some fairly straight forward ways. The following is a list of features which can be easily incorporated into the present framework.

- **Acquiring a Large Block:** Our approach can be easily modified to handle acquisition, rather than liquidation, of a large block with price impact and transactions costs. The individual's objective function can be changed to be $\mathbf{E}[u(X(T) \cdot P(T))]$, and sales can be forbidden (i.e. $dS < 0$ and $\delta < 0$). The remainder of our methodology applies directly.
- **Portfolio Liquidation:** The model can be extended to handle the liquidation of a portfolio by simply applying our approach to each security in the portfolio independently. For example, the price impact and return process of each security can be different. This of course ignores the correlation structure of returns, which may be a critical determinant of an optimal portfolio liquidation policy.
- **Deterministic Time-Varying λ :** If $\lambda = \tilde{\lambda}(t)$, then our numerical solution applies without any modification other than changing λ to $\tilde{\lambda}(t)$ in all the places where it appears. Theorem 3 of course no longer applies, and this special case would have to be handled numerically.
- **Per Share Costs:** To handle per share costs in our discrete trading framework, we simply modify the revenue function as follows

$$\tilde{R}(\delta, P) \equiv R(\delta, P) - k_1(M + X P)\delta,$$

where k_1 is the percentage per share cost. Again, keeping in mind that over a short time horizon $M + X P$ remains relatively constant this is a reasonable way to specify the per share transactions cost. The remainder of the methodology remains unaffected.

Incorporating these elements into our analysis does not pose any new conceptual or computational challenges.

8 Conclusion

Our model suggests a wide range of interesting and potentially fruitful extensions. For example, our present framework does not properly handle the case of portfolio liquidation. Important issues which arise are the correlation in security returns, the effects of multiple securities on the optimal portfolio choice, and the possibility that different securities have different price impacts. Furthermore, it is possible to handle the case of temporary price impact via a straightforward extension of our framework. The idea is that a sale of shares today depresses the stock price, but only temporarily, as liquidity comes back to the market over time. A related extension is the case where the price impact term λ is a stochastic variable.

Hopefully, our solution methodology will allow the effects and determinants of price impact to be studied in a continuous time equilibrium setting. For the most part, previous equilibrium work on trading in the presence of price impact was done in a discrete time setting because the solution of the control problem of an optimizing, risk-averse agent in continuous time was not understood. The analytical benefits of treating the problem in a continuous time setting are substantial. For example, the extant equilibrium discrete time literature has focused exclusively either on agents with constant absolute risk aversion, or on risk-neutral agents. Our analysis allows for the treatment of a much richer, and more realistic, class of utility functions.

Finally, the techniques outlined in this paper should allow for the pricing of derivative securities in the presence of market frictions, such as transaction costs, price impacts, or discrete trading. Even though replication arguments fail to hold in many such settings, we can still price these securities by using the shadow price technique outlined in this paper.

9 Appendix

9.1 Proofs of Theorems 1 and 2

For brevity, we omit the proofs of these Theorems. In both cases, we are dealing with standard singular and impulse control problems respectively. For a detailed treatment of singular control problems, see Shreve and Soner (1994), and for impulse control problems see Eastham and Hastings (1988). Also Hindy, Huang, and Zhu (1997) provide a detailed analytical treatment of a control problem very similar to our own.

9.2 Numerical Method for Continuous Trading Case

In this section we discuss how to numerically solve the free-boundary PDE (29) subject to boundary conditions (25,26,27), and (28). We discretize the state space into a $N_T \times N_\alpha \times N_X$ grid. Given initial values of X_{max} , α_{min} , α_{max} and T , we will have $dx = X_{max}/(N_X - 1)$, $d\alpha = (\alpha_{max} - \alpha_{min})/(N_\alpha - 1)$, $dt = T/(N_T - 1)$. Note that $\alpha_{min} = -\infty$, which is approximated by choosing a sufficiently negative, but finite, number.

It is easy to verify that for the case of continuous trading, the boundary condition (25) implies that condition (28) is exactly zero (i.e. $\gamma Ag - (\lambda + A)g_\alpha - g_X = 0$ where $A \equiv \exp(\alpha)$). The numerical scheme is better behaved if, instead of imposing condition (25), condition (28) is imposed at time T . Discretizing this condition (suppressing the T) yields

$$\gamma Ag(\alpha, X) - (\lambda + A) \frac{g(\alpha, X) - g(\alpha - d\alpha, X)}{d\alpha} - \frac{g(\alpha, X) - g(\alpha, X - dX)}{dX} = 0.$$

Hence for time T we have

$$g(\alpha, X) = \frac{g(\alpha - d\alpha, X)(\lambda + A)dX + g(\alpha, X - dX)d\alpha}{d\alpha + (\lambda + A)dX + \gamma A d\alpha dX} \quad (37)$$

subject to the boundary conditions that

$$g(-\infty, X) = g(\alpha, 0) = 1/\gamma.$$

Given these conditions at time T , we apply the differential operator in (29) to obtain that for time $T - dt$ we have

$$g(\alpha, X, t - dt) = g(\alpha, X, t) + dt \left(r\gamma g(\alpha, X, t) + (\mu - r - \frac{1}{2}\sigma^2)g_\alpha + \frac{1}{2}\sigma^2 g_{\alpha\alpha} \right),$$

where

$$\begin{aligned} g_\alpha &\equiv \frac{g(\alpha + d\alpha, X, t) - g(\alpha - d\alpha, X, t)}{2 d\alpha}, \\ g_{\alpha\alpha} &\equiv \frac{g(\alpha + d\alpha, X, t) - 2g(\alpha, X, t) + g(\alpha - d\alpha, X, t)}{d\alpha^2}. \end{aligned}$$

Note two things: (1) for $\alpha = \alpha_{min}$ we need to impose that $g(\alpha_{min}, X, t) = 1/\gamma \exp(r(T-t)\gamma)$, which for numerical stability is best accomplished by setting $g_\alpha, g_{\alpha\alpha}$ to 0 and applying the procedure above, and (2) and for $X = 0$, the time T boundary conditions insure that $g_\alpha = g_{\alpha\alpha} = 0$.

We now need to check where the free boundary occurs at time $t - dt$. We know that for $X = 0$ and $\alpha = \alpha_{min}$ we are in the no-trade region. The classification step for other points is described below. The ordering of how points are classified is important and must proceed as follows: we start at point $\{\alpha_{min} + d\alpha, dX\}$ and apply the classification step, then move to point $\{\alpha_{min} + 2 d\alpha, dX\}$,

and so on until $\{\alpha_{max}, dX\}$. Then we move on to point $\{\alpha_{min}, 2dX\}$ and repeat the procedure until all the points associated with $2dX$ have been classified. We repeat this until all points associated with X_{max} (i.e. of the form $\{\alpha, X_{max}\}$) have been classified.

The reason that classification must proceed in this manner is that for point $\{\alpha, X\}$ the classification condition (38) relies on points $\{\alpha - d\alpha, X\}$ and $\{\alpha, X - dX\}$, and these need to have already been classified. For each $\{\alpha, X\}$ we need to check whether (the $t - dt$ is suppressed)

$$\gamma Ag(\alpha, X) - (\lambda + A) \frac{g(\alpha, X) - g(\alpha - d\alpha, X)}{d\alpha} - \frac{g(\alpha, X) - g(\alpha, X - dX)}{dX} \quad (38)$$

is negative. As long as this is negative, we are in the no-trade region. As soon as this condition is non-negative for a given α , we classify the point $\{X, \tilde{\alpha}\}$ as being in the trade region. Furthermore, for all points in the trade region condition (38) must be equal to 0, and hence we solve for $g(\tilde{\alpha}, X, t - dt)$ as in (37) above. After all the α points for a given X have been classified in increasing order (i.e. $d\alpha, 2d\alpha, 3d\alpha, \dots, \alpha_{max}$) we move to $X + dX$, and repeat this classification procedure.

Then the differential operator is applied, and we repeat the classification procedure for times $t - 2dt, \dots, 0$.

9.2.1 Numerical Method for No Price Impact Case

This is a straightforward specialization of the method from the previous section.

9.3 Numerical Method for Impulse Trading Case

This methodology is similar to the continuous trading case discussed above. We consider a revenue function $R(\delta, P) = \delta P \exp(-\lambda \delta) - \kappa(M + X P)$. After trade occurs the new price is given by $P \exp(-\lambda \delta)$, although it is clear that other price impact functions could be used instead. Using the state variables α and X , the post trade money amount \tilde{M} can be written as $\tilde{M} = M + R(\delta, P) = M \times c_1$ where

$$c_1 = 1 + \delta e^{\alpha - \lambda \delta} - \kappa(1 + X e^{\alpha}).$$

The time T values are set as follows: $g(\alpha_{min}, X) = g(\alpha, 0) = 1/\gamma$, and

$$g(\alpha, X, T) = \begin{cases} \frac{1}{\gamma} c_1 & \text{if } c_1 \geq 1, \\ \frac{1}{\gamma} & \text{if } c_1 < 1. \end{cases}$$

The time step from t to $t - dt$ is done in the same way as in the section on continuous trading. Let us refer to the γ 's obtained after application of the differential operator from time t to $t - dt$ as \tilde{g} . We now need to insure that condition (19) holds. We note first that

$$\mathbf{E}_{t-dt} J(M(t), X(t), P(t), t) = M^\gamma \times \tilde{g}(\alpha, X, t - dt).$$

Furthermore, we know that all points $\{\alpha_{min}, X\}$ and $\{\alpha, 0\}$ are in the no-trade region.

To classify all other points we proceed as follows. For a point $\{\alpha, X\}$, we need to compute $\mathbf{T}[J]$. Note that for a trade of δ we have

$$\begin{aligned} \alpha &= \tilde{\alpha} - \lambda \delta - \log(c_1), \\ X &= \tilde{X} - \delta, \\ J(\delta) &= M^\gamma c_1^\gamma \tilde{g}(\alpha, X). \end{aligned}$$

where $\tilde{\cdot}$ refers to the pre-trade quantities. We look for the δ which maximizes $J(\delta)$, while restricting our numerical search for δ to the grid, i.e. $\delta = N \times dx$ for some integer N . However, α as defined

above is not necessarily on the grid. Since X is on the grid, we approximate $g(\alpha, X)$ using a 2nd order Taylor expansion as follows. Find α_1 on the grid such that $\alpha \in [\alpha_1, \alpha_1 + d\alpha]$, and then use

$$g(\alpha, X) \approx g(\alpha_1, X) + g_\alpha(\alpha_1, X)(\alpha - \alpha_1) + \frac{1}{2}g_{\alpha\alpha}(\alpha_1, X)(\alpha - \alpha_1)^2,$$

where

$$\begin{aligned} g_\alpha(\alpha_1, X) &\equiv \frac{g(\alpha_1 + d\alpha, X) - g(\alpha_1 - d\alpha, X)}{2d\alpha}, \\ g_{\alpha\alpha}(\alpha_1, X) &\equiv \frac{g(\alpha_1 + d\alpha, X) - 2g(\alpha_1, X) + g(\alpha_1 - d\alpha, X)}{d\alpha^2}. \end{aligned}$$

For $\alpha_1 = \alpha_{min}$, we use

$$\begin{aligned} g_\alpha(\alpha_{min}, X) &\equiv \frac{g(\alpha_{min} + d\alpha, X) - g(\alpha_{min}, X)}{d\alpha}, \\ g_{\alpha\alpha}(\alpha_{min}, X) &\equiv 0. \end{aligned}$$

From an implementation point of view, the 2nd order approximation works far better than the first order approximation.

Let us now refer to the post-classification values of g without the $\tilde{\cdot}$. Once $\mathbf{T}J \equiv \sup_\delta J(\delta)$ has been obtained, we check to see if

$$\tilde{g}(\tilde{\alpha}, \tilde{X}) \leq c_1^\gamma \tilde{g}(\alpha, X),$$

and if this holds then the pair $\{\tilde{\alpha}, \tilde{X}\}$ is in the trade region, and we set $g(\tilde{\alpha}, \tilde{X}) = c_1^\gamma \tilde{g}(\alpha, X)$. Otherwise the pair is in the no-trade region and $g(\tilde{\alpha}, \tilde{X}) = \tilde{g}(\tilde{\alpha}, \tilde{X})$.

Once the classification step has been performed for $t - dt$, we repeat the application of the differential operator and the classification step for times $t - 2dt, \dots, 0$.

References

- Almgren, R. and Chriss, N., 1998, "Optimal liquidation," working paper, University of Chicago.
- Bertsimas, D. and Lo, A.W., 1998, "Optimal control of execution costs," *Journal of Financial Markets*, 1, 1–50.
- Chan, L.K.C. and J. Lakonishok, 1995, "The behavior of stock prices around institutional trades," *Journal of Finance*, 50, 1147–1174.
- Eastham, J.F. and K.J. Hastings, 1988, "Optimal Impulse Control of Portfolios," *Mathematics of Operations Research* 13 (4), 588-605.
- Hindy, A., C.-f. Huang, and S.H. Zhu, 1997, "Numerical analysis of a free boundary singular control problem in financial economics," *Journal of Economic Dynamics and Control*, 21, 297–327.
- Holthausen, R., R. Leftwich, and D. Mayers, 1990, "Large block transactions, the speed of response of temporary and permanent stock-price effects," *Journal of Financial Economics*, 26, 71–95.
- Keim, D. and A. Madhavan, 1996, "The upstairs market for large-block transactions: Analysis and measurement of price effects," *Review of Financial Studies*, 9, 1–36.
- Kyle, A.S., 1985, "Continuous auctions and insider trading," *Econometrica*, 53 (6), 1315–1335.
- Liu, H. and J. Yong, 2001, "Option pricing in an illiquid market," mimeo, Olin School of Business, Washington University.
- Merton, R.C., 1971, "Optimum consumption and portfolio rules in a continuous-time model," *Journal of Economic Theory*, 3, 373–413.
- Morton, A.J. and Pliska, S.R., 1995, "Optimal Portfolio Management with Fixed Transaction Costs," *Mathematical Finance*, 5, 337-356.
- Schroeder, M., 1998, "Optimal Portfolio Selection with Fixed Transaction Costs," working paper, Northwestern University.
- Shreve, S. and H. Soner, 1994, "Optimal investment and consumption with transaction costs," *Annals of Applied Probability*, 4, 609–692.
- Stanzl, W. and G. Huberman, 2000, "Optimal liquidity trading," mimeo, Yale School of Management.
- Subramanian, A. and Jarrow, R.A., 2001, "The Liquidity Discount," *Mathematical Finance*, 11 (4).
- Vayanos, D., 2000, "Strategic Trading in a Dynamic Noisy Market," forthcoming in *Journal of Finance*.

Table 1: Results of 10,000 simulation runs for liquidating a large block of shares. The table reports means, standard deviations, and percentiles over the simulation runs. The variables marked with \cdot^* refer to the optimal liquidation policy, and the unmarked variables refer to the naive policy of linear liquidation. $R(T)$ refers to the simple return of each strategy at time T , and is defined by $R_T \equiv M(T)/(M(0^-) + X(0^-) \cdot P(0^-)) - 1$. The variable $\Pi(T)$ is the share weighted execution price as of time T , and is defined as $(\sum_{s=0, \dots, T} \delta(s) \cdot P(s))/(P(0^-) \cdot \sum_s \delta(s))$, where $\delta(s)$ is the number of shares sold at time s . $u(M(T))$ is the utility over time T money. The initial pre-trade values for the state variables are $M(0^-) = \exp(-2)$, $P(0^-) = 1$, $X(0^-) = 10$, and the number of trading periods is $N_T = 20$. The parameter values are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$, $\gamma = -3$, $k = 0$.

Variable	Mean	SD	1.0	2.5	5.0	50.0	95.0	97.5	99.0
$R^*(T)$	-0.04229	0.04255	-0.1352	-0.1226	-0.1094	-0.0436	0.0292	0.0445	0.0632
$R(T)$	-0.04123	0.05055	-0.1505	-0.1358	-0.1206	-0.0427	0.0445	0.0621	0.0844
$\Pi^*(T)$	0.95424	0.04303	0.8603	0.8730	0.8864	0.9529	1.0266	1.0420	1.0609
$\Pi(T)$	0.95571	0.05114	0.8451	0.8601	0.8754	0.9543	1.0424	1.0602	1.0828

Variable	Mean	S.E.	$J(0^-)$
$u(M^*(T))$	-0.00036879	0.00000049	-0.00036864
$u(M(T))$	-0.00036934	0.00000058	—

Table 2: The result of a single simulation run. The variables marked with $*$ refer to the optimal strategy, and the unmarked variables refer to the naive strategy of linear liquidation. The variables P, X, δ, M represent, respectively, the price, shares still held, trade amount, and money after trade has taken place in a given period. A period n is mapped to time as follows $n \rightarrow T \times (n-1)/(N_T-1)$. The initial pre-trade values for the state variables are $M(0^-) = \exp(-2)$, $P(0^-) = 1$, $X(0^-) = 10$, and the number of trading periods is $N_T = 20$. The parameter values are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$, $\gamma = -3$, $k = 0$.

Period	P^*	X^*	δ^*	M^*	P	X	δ	M
1	0.9903	9.0301	0.9699	1.0959	0.9950	9.5000	0.5000	0.6328
2	0.9907	8.1635	0.8666	1.9547	0.9990	9.0000	0.5000	1.1325
3	0.9848	7.4077	0.7557	2.6995	0.9956	8.5000	0.5000	1.6306
4	0.9720	6.7389	0.6689	3.3503	0.9843	8.0000	0.5000	2.1232
5	0.9458	6.1368	0.6020	3.9206	0.9588	7.5000	0.5000	2.6032
6	0.9405	5.6017	0.5351	4.4249	0.9538	7.0000	0.5000	3.0807
7	0.9347	5.1001	0.5017	4.8950	0.9479	6.5000	0.5000	3.5555
8	0.9102	4.6318	0.4682	5.3225	0.9227	6.0000	0.5000	4.0178
9	0.9022	4.2267	0.4051	5.6893	0.9137	5.5000	0.5000	4.4757
10	0.9167	3.8521	0.3746	6.0342	0.9273	5.0000	0.5000	4.9405
11	0.9363	3.4842	0.3679	6.3803	0.9459	4.5000	0.5000	5.4147
12	0.9987	3.1163	0.3679	6.7494	1.0076	4.0000	0.5000	5.9199
13	0.9773	2.7484	0.3679	7.1107	0.9847	3.5000	0.5000	6.4138
14	0.9850	2.4080	0.3404	7.4479	0.9908	3.0000	0.5000	6.9109
15	0.9901	2.0401	0.3679	7.8141	0.9946	2.5000	0.5000	7.4101
16	0.9973	1.6722	0.3679	8.1830	1.0006	2.0000	0.5000	7.9123
17	0.9542	1.3043	0.3679	8.5362	0.9561	1.5000	0.5000	8.3924
18	0.9395	0.9030	0.4013	8.9155	0.9404	1.0000	0.5000	8.8648
19	0.9453	0.4682	0.4348	9.3289	0.9456	0.5000	0.5000	9.3400
20	0.9527	-0.0000	0.4682	9.7774	0.9527	0.0000	0.5000	9.8188

Table 3: Results of 10,000 simulation runs for liquidating a large block of shares. The table reports means, standard deviations, and percentiles over the simulation runs. The variables marked with \cdot^* refer to the optimal liquidation policy, and the unmarked variables refer to the naive policy of linear liquidation. $R(T)$ refers to the simple return of each strategy at time T , and is defined by $R_T \equiv M(T)/(M(0^-) + X(0^-) \cdot P(0^-)) - 1$. The variable $\Pi(T)$ is the share weighted execution price as of time T , and is defined as $(\sum_{s=0, \dots, T} \delta(s) \cdot P(s))/(P(0^-) \cdot \sum_s \delta(s))$, where $\delta(s)$ is the number of shares sold at time s . $u(M(T))$ is the utility over time T money. The initial pre-trade values for the state variables are $M(0^-) = \exp(-2)$, $P(0^-) = 1$, $X(0^-) = 10$, and the number of trading periods is $N_T = 20$. The parameter values are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$, $\gamma = 1$, $k = 0$.

Variable	Mean	SD	1.0	2.5	5.0	50.0	95.0	97.5	99.0
$R^*(T)$	-0.03965	0.06436	-0.1770	-0.1580	-0.1420	-0.0412	0.0703	0.0927	0.1185
$R(T)$	-0.04052	0.05025	-0.1485	-0.1340	-0.1207	-0.0416	0.0446	0.0619	0.0848
$\Pi^*(T)$	0.95814	0.06515	0.8191	0.8383	0.8545	0.9565	1.0694	1.0921	1.1182
$\Pi(T)$	0.95644	0.05084	0.8472	0.8619	0.8753	0.9553	1.0426	1.0600	1.0832

Variable	Mean	S.E.	$J(0^-)$
$u(M^*(T))$	9.73348814	0.00652356	9.72609292
$u(M(T))$	9.72468597	0.00509316	—

Table 4: The result of a single simulation run. The variables marked with * refer to the optimal strategy, and the unmarked variables refer to the naive strategy of linear liquidation. The variables P, X, δ, M represent, respectively, the price, shares still held, trade amount, and money after trade has taken place in a given period. A period n is mapped to time as follows $n \rightarrow T \times (n-1)/(N_T-1)$. The initial pre-trade values for the state variables are $M(0^-) = \exp(-2), P(0^-) = 1, X(0^-) = 10$, and the number of trading periods is $N_T = 20$. The parameter values are $\mu = 0.14, \sigma = 0.3, \lambda = 0.01, r = 0.05, T = 0.1, N_T = 20, \gamma = 1, k = 0$.

Period	P^*	X^*	δ^*	M^*	P	X	δ	M
1	0.9997	9.9666	0.0334	0.1688	0.9950	9.5000	0.5000	0.6328
2	1.0084	9.8887	0.0778	0.2473	0.9995	9.0000	0.5000	1.1328
3	0.9942	9.7549	0.1338	0.3804	0.9818	8.5000	0.5000	1.6240
4	0.9968	9.5621	0.1928	0.5727	0.9814	8.0000	0.5000	2.1151
5	0.9902	9.3280	0.2341	0.8046	0.9722	7.5000	0.5000	2.6018
6	1.0182	9.0605	0.2676	1.0773	0.9975	7.0000	0.5000	3.1012
7	1.0113	8.7542	0.3063	1.3873	0.9888	6.5000	0.5000	3.5964
8	1.0402	8.3764	0.3778	1.7807	1.0158	6.0000	0.5000	4.1052
9	1.0742	7.9750	0.4013	2.2123	1.0480	5.5000	0.5000	4.6303
10	1.0683	7.4734	0.5017	2.7488	1.0422	5.0000	0.5000	5.1526
11	1.0415	6.9382	0.5351	3.3068	1.0164	4.5000	0.5000	5.6622
12	1.0190	6.3697	0.5686	3.8871	0.9952	4.0000	0.5000	6.1612
13	1.0354	5.7342	0.6355	4.5460	1.0125	3.5000	0.5000	6.6691
14	1.0580	5.0521	0.6821	5.2689	1.0365	3.0000	0.5000	7.1891
15	1.0484	4.3498	0.7023	6.0066	1.0291	2.5000	0.5000	7.7056
16	1.0473	3.5806	0.7692	6.8139	1.0309	2.0000	0.5000	8.2231
17	1.0768	2.7456	0.8349	7.7147	1.0634	1.5000	0.5000	8.7570
18	1.1022	1.8761	0.8696	8.6751	1.0926	1.0000	0.5000	9.3056
19	1.1506	0.9699	0.9062	9.7201	1.1452	0.5000	0.5000	9.8806
20	1.0997	0.0000	0.9699	10.7892	1.0997	0.0000	0.5000	10.4331

Table 5: Results of 10,000 simulation runs for liquidating a large block of shares. The table reports means, standard deviations, and percentiles over the simulation runs. The variables marked with \cdot^* refer to the optimal liquidation policy, and the unmarked variables refer to the naive policy of linear liquidation. $R(T)$ refers to the simple return of each strategy at time T , and is defined by $R_T \equiv M(T)/(M(0^-) + X(0^-) \cdot P(0^-)) - 1$. The variable $\Pi(T)$ is the share weighted execution price as of time T , and is defined as $(\sum_{s=0, \dots, T} \delta(s) \cdot P(s))/(P(0^-) \cdot \sum_s \delta(s))$, where $\delta(s)$ is the number of shares sold at time s . $u(M(T))$ is the utility over time T money. The initial pre-trade values for the state variables are $M(0^-) = \exp(-2)$, $P(0^-) = 1$, $X(0^-) = 10$, and the number of trading periods is $N_T = 20$. The parameter values are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$, $\gamma = -3$, $k = 0.001$.

Variable	Mean	SD	1.0	2.5	5.0	50.0	95.0	97.5	99.0
$R^*(T)$	-0.05369	0.03278	-0.1253	-0.1149	-0.1056	-0.0547	0.0027	0.0137	0.0272
$R(T)$	-0.06036	0.04982	-0.1687	-0.1531	-0.1389	-0.0625	0.0250	0.0445	0.0643
$\Pi^*(T)$	0.94922	0.03330	0.8765	0.8870	0.8965	0.9482	1.0065	1.0177	1.0314
$\Pi(T)$	0.95592	0.05115	0.8446	0.8607	0.8754	0.9536	1.0436	1.0635	1.0838

Variable	Mean	S.E.	$J(0^-)$
$u(M^*(T))$	-0.00038052	0.00000039	-0.00038055
$u(M(T))$	-0.00039242	0.00000062	—

Table 6: The result of a single simulation run. The variables marked with * refer to the optimal strategy, and the unmarked variables refer to the naive strategy of linear liquidation. The variables P, X, δ, M represent, respectively, the price, shares still held, trade amount, and money after trade has taken place in a given period. A period n is mapped to time as follows $n \rightarrow T \times (n-1)/(N_T-1)$. The initial pre-trade values for the state variables are $M(0^-) = \exp(-2), P(0^-) = 1, X(0^-) = 10$, and the number of trading periods is $N_T = 20$. The parameter values are $\mu = 0.14, \sigma = 0.3, \lambda = 0.01, r = 0.05, T = 0.1, N_T = 20, \gamma = -3, k = 0.001$.

Period	P^*	X^*	δ^*	M^*	P	X	δ	M
1	0.9844	8.4281	1.5719	1.6726	0.9950	9.5000	0.5000	0.6227
2	1.0121	6.8913	1.5368	3.2181	1.0336	9.0000	0.5000	1.1292
3	0.9892	5.4531	1.4381	4.6313	1.0198	8.5000	0.5000	1.6290
4	0.9497	4.0805	1.3727	5.9262	0.9876	8.0000	0.5000	2.1132
5	0.9047	2.7092	1.3712	7.1587	0.9491	7.5000	0.5000	2.5786
6	0.8600	2.7092	0.0000	7.1606	0.8977	7.0000	0.5000	3.0187
7	0.8438	2.7092	0.0000	7.1625	0.8764	6.5000	0.5000	3.4486
8	0.8227	2.7092	0.0000	7.1644	0.8502	6.0000	0.5000	3.8656
9	0.8218	2.7092	0.0000	7.1663	0.8451	5.5000	0.5000	4.2802
10	0.8410	2.7092	0.0000	7.1682	0.8605	5.0000	0.5000	4.7025
11	0.8395	2.7092	0.0000	7.1700	0.8547	4.5000	0.5000	5.1221
12	0.8774	2.7092	0.0000	7.1719	0.8888	4.0000	0.5000	5.5587
13	0.8791	2.7092	0.0000	7.1738	0.8861	3.5000	0.5000	5.9941
14	0.9117	2.7092	0.0000	7.1757	0.9143	3.0000	0.5000	6.4436
15	0.9269	2.7092	0.0000	7.1776	0.9250	2.5000	0.5000	6.8986
16	0.9292	2.7092	0.0000	7.1795	0.9226	2.0000	0.5000	7.3525
17	0.9109	2.7092	0.0000	7.1814	0.9000	1.5000	0.5000	7.7953
18	0.9259	2.7092	0.0000	7.1833	0.9102	1.0000	0.5000	8.2432
19	0.8906	1.3712	1.3380	8.3671	0.8829	0.5000	0.5000	8.6777
20	0.8895	0.0000	1.3712	9.5794	0.8895	0.0000	0.5000	9.1156

Table 7: Results of 10,000 simulation runs for liquidating a large block of shares. The table reports means, standard deviations, and percentiles over the simulation runs. The variables marked with \cdot^* refer to the optimal liquidation policy, and the unmarked variables refer to the naive policy of linear liquidation. $R(T)$ refers to the simple return of each strategy at time T , and is defined by $R_T \equiv M(T)/(M(0^-) + X(0^-) \cdot P(0^-)) - 1$. The variable $\Pi(T)$ is the share weighted execution price as of time T , and is defined as $(\sum_{s=0, \dots, T} \delta(s) \cdot P(s))/(P(0^-) \cdot \sum_s \delta(s))$, where $\delta(s)$ is the number of shares sold at time s . $u(M(T))$ is the utility over time T money. The initial pre-trade values for the state variables are $M(0^-) = \exp(-2)$, $P(0^-) = 1$, $X(0^-) = 10$, and the number of trading periods is $N_T = 20$. The parameter values are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$, $\gamma = 1$, $k = 0.001$.

Variable	Mean	SD	1.0	2.5	5.0	50.0	95.0	97.5	99.0
$R^*(T)$	-0.04880	0.08161	-0.2229	-0.1984	-0.1772	-0.0527	0.0899	0.1209	0.1603
$R(T)$	-0.06020	0.05037	-0.1698	-0.1530	-0.1390	-0.0619	0.0266	0.0452	0.0656
$\Pi^*(T)$	0.95581	0.08316	0.7784	0.8034	0.8250	0.9518	1.0971	1.1287	1.1689
$\Pi(T)$	0.95608	0.05172	0.8435	0.8608	0.8751	0.9543	1.0451	1.0644	1.0851

Variable	Mean	S.E.	$J(0^-)$
$u(M^*(T))$	9.64074302	0.00827166	9.63545541
$u(M(T))$	9.52521520	0.00510508	—

Table 8: The result of a single simulation run. The variables marked with $*$ refer to the optimal strategy, and the unmarked variables refer to the naive strategy of linear liquidation. The variables P, X, δ, M represent, respectively, the price, shares still held, trade amount, and money after trade has taken place in a given period. A period n is mapped to time as follows $n \rightarrow T \times (n-1)/(N_T-1)$. The initial pre-trade values for the state variables are $M(0^-) = \exp(-2)$, $P(0^-) = 1$, $X(0^-) = 10$, and the number of trading periods is $N_T = 20$. The parameter values are $\mu = 0.14$, $\sigma = 0.3$, $\lambda = 0.01$, $r = 0.05$, $T = 0.1$, $N_T = 20$, $\gamma = 1$, $k = 0.001$.

Period	P^*	X^*	δ^*	M^*	P	X	δ	M
1	1.0000	10.0000	0.0000	0.1353	0.9950	9.5000	0.5000	0.6227
2	1.0326	10.0000	0.0000	0.1354	1.0224	9.0000	0.5000	1.1237
3	1.0770	10.0000	0.0000	0.1354	1.0609	8.5000	0.5000	1.6437
4	1.0870	10.0000	0.0000	0.1354	1.0654	8.0000	0.5000	2.1661
5	1.0969	10.0000	0.0000	0.1355	1.0698	7.5000	0.5000	2.6908
6	1.1023	10.0000	0.0000	0.1355	1.0697	7.0000	0.5000	3.2156
7	1.1468	10.0000	0.0000	0.1355	1.1073	6.5000	0.5000	3.7591
8	1.1632	10.0000	0.0000	0.1356	1.1175	6.0000	0.5000	4.3078
9	1.1951	10.0000	0.0000	0.1356	1.1425	5.5000	0.5000	4.8690
10	1.1811	10.0000	0.0000	0.1357	1.1235	5.0000	0.5000	5.4209
11	1.2450	10.0000	0.0000	0.1357	1.1784	4.5000	0.5000	6.0002
12	1.2687	10.0000	0.0000	0.1357	1.1948	4.0000	0.5000	6.5878
13	1.2760	10.0000	0.0000	0.1358	1.1957	3.5000	0.5000	7.1760
14	1.2535	10.0000	0.0000	0.1358	1.1687	3.0000	0.5000	7.7509
15	1.2407	8.4615	1.5385	2.0319	1.1689	2.5000	0.5000	8.3262
16	1.1972	6.8591	1.6024	3.9386	1.1404	2.0000	0.5000	8.8874
17	1.1522	5.2203	1.6388	5.8159	1.1102	1.5000	0.5000	9.4337
18	1.1282	3.5452	1.6752	7.6955	1.0998	1.0000	0.5000	9.9750
19	1.1224	1.8060	1.7391	9.6378	1.1078	0.5000	0.5000	10.5204
20	1.1422	0.0000	1.8060	11.6915	1.1422	0.0000	0.5000	11.0832