NORTHWESTERN UNIVERSITY

Essays in Finance

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Finance

By

Caroline-Anne Sasseville

EVANSTON, ILLINOIS

December 2007

ABSTRACT

Essays in Finance

Caroline-Anne Sasseville

In chapter 1, we study the investment behavior of firms faced with uncertainty, irreversibility and non-convex adjustment costs when output can be stored. A discrete-time dynamic optimization problem is solved numerically using neural networks to study the investment problem. Our numerical results suggest that whether or not firms can store their output can have a significant impact on the investment decision, particularly when demand is stochastic. This chapter is an important first step to understanding the equilibrium price dynamics of storable commodities. Such an equilibrium model is proposed in chapter 2.

The purpose of chapter 2 is to model the price dynamics of a storable commodity in a production economy framework. Our model consists of an infinite horizon economy with a risk-averse agent and two goods: a liquid asset and a commodity. The primary purpose served by the commodity is endogenously determined by the supply of liquidity. The model is solved numerically and the results suggest that the joint dynamics of the commodity spot price and convenience yield depend on the nature of the commodity. In

particular, we find that when a commodity serves primarily as a store of value, the spot price and convenience yield can move in opposite directions.

The purpose of chapter 3 is to investigate investors' responses to ownership disclosures in a specific context. We examine in real-time the recommendations made by analysts on one of the *Making Money Now* segments, which air on CNBC while markets are open. The recommendations are followed by disclosures, which are mandatory under Title V of the Sarbanes Oxley Act of 2002, and NYSE and NASD rules. We measure the reaction to both the recommendation event and the disclosure event. While we find an immediate response in price, volume and order flow following the recommendation event, there appears to be no response following the disclosure event. Our results should be useful to policymakers who must balance the costs imposed on those who must disclose with the benefits to individual investors who gain from the disclosure.

Acknowledgements

I would like to thank the members of my dissertation committee, Robert McDonald, Janice Eberly, Ernst Schaumburg, and Mark Satterthwaite, for their helpful comments and discussions. I am particularly grateful to Robert McDonald for being an outstanding mentor, for his continuous support, and for his patience and understanding.

I am also grateful to Torben Andersen, Andrea Eisfeldt, Michael Fishman, Andrew Hertzberg, Yael Hochberg, Ravi Jagannathan, Arvind Krishnamurthy, Camelia Kuhnen, Deborah Lucas, David Matsa, Mitchell Petersen, Paola Sapienza, Constantinos Skiadas, Annette Vissing-Jørgensen, and Neng Wang for helpful comments and suggestions. Special thanks to Patricia Ledesma Liébana and Jim Guo for providing exceptional technical support.

Many thanks to my fellow doctoral students, Joey Engelberg, Paul Gao, Grace Koo, Niels Schühle, Georgios Skoulakis, Nicolas Vincent, and Jared Williams for their support and willingness to help. I am especially grateful to Yanik Proulx and Véronique Harmignies for helping me through a very difficult period of my studies. Special thanks to Jeremiah Heller for saving me from homemade lattes.

Last, but certainly not least, I would like to thank my family for their unconditional love and support.

À mes parents,

Danièle et Yves

Table of Contents

List of Tables

List of Figures

CHAPTER 1

Investment Under Uncertainty: Q-theory and the Possibility of Storage

1.1. Introduction

Tobin (1969) argued that the investment of a firm should be an increasing function of the ratio of the market value of the firm's assets to the replacement cost of capital (this ratio became known as Tobin's q). Among others, Mussa (1977) linked the adjustment cost and the q literatures by showing that, in a model with convex adjustment costs, what determines the optimal rate of investment is the marginal benefit to the firm of an additional unit of capital (a measure known as marginal q) since the optimal rate of investment is the rate at which marginal q equals the marginal cost of adjustment. Hayashi (1982) showed that when a firm's profit function is homogeneous of degree one in capital and the adjustment cost function is convex and homogeneous of degree one in capital and investment, average q and marginal q are equal. Tobin's q is therefore the relevant measure to determine the optimal investment rate if the conditions specified by Hayashi are met.

Despite the strong theoretical support for q-theory, subsequent empirical research showed little support for Tobin's q as a sufficient statistic for investment. Investment regressions using q as an explanatory variable had very low R^2's and other variables, such as cash flow and profits, were found to have more explanatory power than q. The poor empirical performance of q suggested that perhaps the assumptions of models with convex adjustment costs were too stringent and models which allow for more frictions should be considered.

Abel and Eberly (1994) propose a model of irreversible investment with non-convex adjustment costs. The non-convexity in the adjustment cost function comes from the presence of a fixed cost associated with investment. The presence of a fixed component in the adjustment cost function is empirically supported. For instance, using U.S. Census Bureau micro data, Doms and Dunne (1998) study the investment patterns of 13,700 manufacturing plants. They find evidence of lumpy investment, consistent with the presence of a fixed cost.

The main finding of Abel and Eberly (1994) is that while investment is a non-decreasing function of q, it is not a monotonically increasing function of q. More precisely, they find that the introduction of a fixed cost leads to a region where investment is insensitive to q, which leads to three investment regimes (disinvestment, no-investment, investment). Using Compustat data for manufacturing firms, Barnett and Sakellaris (1998) tested this finding of different investment regimes and found support for Abel and Eberly's finding of a non-decreasing response of investment to q. Using Compustat data to study the importance of nonlinearities in the relationship between investment and q which may arise due to fixed costs of adjustment, Abel and Eberly (2002) also find support for regimes of different sensitivities to q.

A common assumption in the literature merging Tobin's q and adjustment costs is that the output good cannot be stored. Allowing for the possibility that the output can be stored could reduce the effect of a fixed cost of adjustment. In the continuous-time framework of Abel and Eberly (1994), the effect of storage on the optimality condition for investment is not readily apparent since the effect of investment on the marginal benefit of storage is a second-order effect. However, firms which can store their output may benefit not only from the increased level of output produced, but also from the additional flexibility to respond to unexpected future demand and production shocks without changing the production process. The benefits of storage may in turn affect the size of the no-investment region.

Motivated by this idea that allowing for the possibility of storage could dampen the effect of adjustment costs, we propose a model of investment with non-convex adjustment costs, irreversible investment and the possibility of storage. The theoretical framework is as in Abel and Eberly (1994), except that output can now be stored. Because the effect of storage on the optimal investment decision is not readily apparent when we solve the model analytically, we solve numerically a discrete-time dynamic optimization problem to gain more insight on the effect of adding the possibility of storage on the investment decision. We find that whether or not a firm can store its output does affect its investment behavior. As such, our work can be considered to be a first step towards an equilibrium model for the price of storable commodities.

The purpose of this chapter is to study the investment behavior of firms which can store their output. The structure of the paper is as follows. In section 2, we present the infinite horizon, continuous-time model and derive the first-order conditions necessary for the optimality of investment when the output can be stored. In section 3, we present the discrete-time model and describe the approach used to solve it numerically. In section 4, we present our numerical results. Finally, section 5 concludes.

1.2. Continuous-time investment model with storage

The framework here is similar to that in Abel and Eberly (1994). Let us consider a firm that chooses the investment rate (represented by I_t) to produce an output, which can be stored and sold in future periods if it is not sold in the current period. Let $\pi(S_t, P_t, Q_t, K_t)$ be the rate of the profit flow (such that the profit over an interval of length dt is $\pi(S_t, P_t, Q_t, K_t)$), where S_t represents the quantity sold at time t, P_t represents a random component in the profit flow at time t, Q_t represents the inventory level at time t, and K_t represents the capital stock at time t. Let the total cost of investment be represented by $a(I_t, K_t)1_{\{I_t \neq 0\}}$, where $1_{\{I_t \neq 0\}}$ is equal to 1 if the firm decides to invest. The function $a(I_t, K_t)$ is the function referred to as the "adjustment cost function"by Abel and Eberly. This function has three components: one to represent the purchase/sale of capital, one to represent the quadratic adjustment costs and one to represent the fixed cost (note that the fixed cost component is equal to $a(0, K_t)$). Assume that $a(I_t, K_t)$ is twice continuously differentiable (except possibly at 0), and that $a_I(I_t, K_t) > 0$ and that $a_{II}(I_t, K_t) > 0$ for all $t > 0$. Finally, assume that $\pi(S_t, P_t, Q_t, K_t)$ is twice continuously differentiable, and that $\pi_S(S_t, P_t, Q_t, K_t) > 0$ and that $\pi_{SS}(S_t, P_t, Q_t, K_t) < 0$ for all $t > 0$.

At time t, for all $s > 0$, the firm faces the following maximization problem:

$$V(P_t, Q_t, K_t) = \max_{\{S_{t+s}, I_{t+s}\}} \int_0^\infty e^{-rs} E_t \begin{pmatrix} \pi(S_{t+s}, P_{t+s}, Q_{t+s}, K_{t+s}) \\ -a(I_{t+s}, K_{t+s})1_{\{I_{t+s} \neq 0\}} \end{pmatrix} ds$$

s.t.

$$S_{t+s} \leq Q_{t+s} + f(K_{t+s}) \tag{1.1}$$

$$S_{t+s} \geq 0 \tag{1.2}$$

where r is the discount rate. The first inequality constraint (equation 1.1) represents the fact the firm cannot sell more than the quantity on hand. This constraint is referred to as the "inventory constraint". Solving this maximization problem using the Bellman equation, we have

$$\begin{aligned} rV(P_t, Q_t, K_t) &= \max_{\{S_t, I_t\}} \pi(S_t, P_t, Q_t, K_t) - a(I_t, K_t) 1_{\{I_t \neq 0\}} \\ &\quad + \frac{1}{dt} E_t \left[dV\right] - \lambda_1 \left(Q_t + f(K_t) - S_t\right) - \lambda_2 S_t \end{aligned}$$

where λ_1 and λ_2 are Lagrange multipliers representing the inventory and the non-negativity constraints (equations 1.1 and 1.2), respectively. Now, to solve this problem further, we must specify the dynamics for the state variables P_t, Q_t and K_t. Let us assume that the random component P_t evolves as according to the diffusion process:

$$dP_t = \mu\left(P_t\right) dt + \sigma_P \left(P_t\right) dW_P$$

where W_P is a standard Wiener process. For K_t, we have that, in every period, the capital stock is equal to the undepreciated capital stock from the previous period plus any investment made at time t. So we have

$$dK_t = (I_t - \delta K_t) dt$$

where δ represents the depreciation rate of capital. For the inventory, we have that the level of the inventory is equal to the inventory carried over from last period, plus the difference between the current period's production and sales. Put differently,

$$dQ_t = (f(K_t) - S_t)dt.$$

It is important to note that I_t does not directly enter the equation for the evolution of the inventory since investment only has a second order effect on the level of the inventory.

Using Ito's Lemma (and dropping time subscripts), we get

$$dV = \left(V_Q dQ + V_K dK + V_P dP + \frac{1}{2} V_{PP} \sigma_P^2 P^2 dt\right).$$

So we have the following maximization problem:

$$\begin{aligned} rV &= \max_{\{S,I\}} \pi(S,P,Q,K) - a(I,K) 1_{\{I \neq 0\}} + V_Q(f(K) - S) \\ &\quad + V_K(I - \delta K) + \mu(P) V_P + \frac{1}{2} V_{PP} \sigma_P^2 (P) \\ &\quad - \lambda_1 (Q + f(K) - S) - \lambda_2 S. \end{aligned}$$

Assuming for now that $I \neq 0$, we have the following first order conditions:

$$\pi_S(S,P,Q,K) - V_Q + \lambda_1 - \lambda_2 = 0 \tag{1.3}$$

$$V_K - a_I(I,K) = 0. \tag{1.4}$$

These first order conditions are intuitively appealing. From (1.3), we have that the optimal quantity to sell depends on the trade-off between the marginal benefit of selling and the marginal benefit of storing an additional unit. From (1.4), we have that the optimal investment rate is the rate at which the marginal benefit of one unit of installed capital equals the marginal cost of installing one additional unit of capital.

From (1.3) and the above constraint, we have that S^* solves

$$\pi_S\left(S^*(K,P,V_Q),P,Q,K\right) = V_Q \quad \text{if} \quad \pi_S\left(Q+f(K),P,Q,K\right) < V_Q < \pi_S\left(0,P,Q,K\right)$$

and

$$S^*(K,P,V_Q,K) = \begin{cases} 0 & \text{if} \quad V_Q \geq \pi_S\left(0,P,Q,K\right) \\ Q+f(K) & \text{if} \quad V_Q \leq \pi_S\left(Q+f(K),P,Q,K\right). \end{cases}$$

From (1.4) and from the conditions imposed on the function $a(I,K)$, we have

$$V_K = a_I(I^*(V_K,K),K) \quad \text{if} \quad V_K > a_I(0,K)^+ \quad \text{or} \quad V_K < a_I(0,K)^-. \tag{1.5}$$

Otherwise,

$$I^*(V_K,K) = 0$$

where $a_I(0,K)^+$ and $a_I(0,K)^-$ represents the right-hand and left-hand derivatives of $a(I,K)$ with respect to I, evaluated at $I=0$, respectively. Since (1.5) was obtained by assuming that $1_{\{I\neq 0\}} = 1$, we now need to determine when this will actually be the case. From (1.3), we have that the only terms which affect the optimal level of I are

$$V_K I - a(I,K)1_{\{I\neq 0\}}.$$

Let $\varphi(V_K,K) = \max_I V_K I - a(I,K) = V_K I^*(V_K,K) - a(I^*(V_K,K),K)$. Since $\varphi(V_K,K)$ can be shown to be a convex function which attains its minimum value when $I^* = 0$, we know that $\varphi(V_K, K)$ attains its minimum value when $a_I(0,K)^- < V_K < a_I(0,K)^+$. This means that we can only have $\varphi(V_K, K) > 0$ for values of V_K satisfying either $V_K > a_I(0,K)^+$ or $V_K < a_I(0,K)^-$. In particular, we have that $\varphi(V_K, K) > 0$ if $V_K < V_{K_1}$ or $V_K > V_{K_2}$ where V_{K_1} and V_{K_2} are the smallest and largest roots of $\varphi(V_K, K)$, respectively. Since $V_{K_1} < a_I(0,K)^-$ and $V_{K_2} > a_I(0,K)^+$, we have

$$V_K = a_I(I^*(V_K, K), K) \quad \text{if} \quad V_K > V_{K_2} \quad \text{or} \quad V_K < V_{K_1}.$$

Otherwise,

$$I^*(V_K, K) = 0.$$

Figure 1.1 summarizes the above discussion. This figure illustrates the effect of the irreversibility of investment and of the fixed cost on the no-investment region. As shown in this figure, irreversibility widens the no-investment region, but is not necessary for an inaction region to exist.

1.3. Discrete-time problem

In the previous section, we derived optimality conditions for the investment rate and the quantity sold in the continuous-time framework. From the first order conditions, it seems that the optimality condition for investment is independent of the marginal benefit of storage. Firms may, nonetheless benefit from holding inventories. For instance, inventories can give producers the flexibility to meet to unexpected demand shocks without having to adjust the production process. This means that the marginal benefit of an additional unit of capital in place may be higher for firms which can store their output. In the context of q-theory with non-convex adjustment costs, this suggests that firms

Figure 1.1: Investment region with fixed costs and irreversible investment. This figure illustrates the effect of the fixed cost and of the irreversibility on the no-investment region.

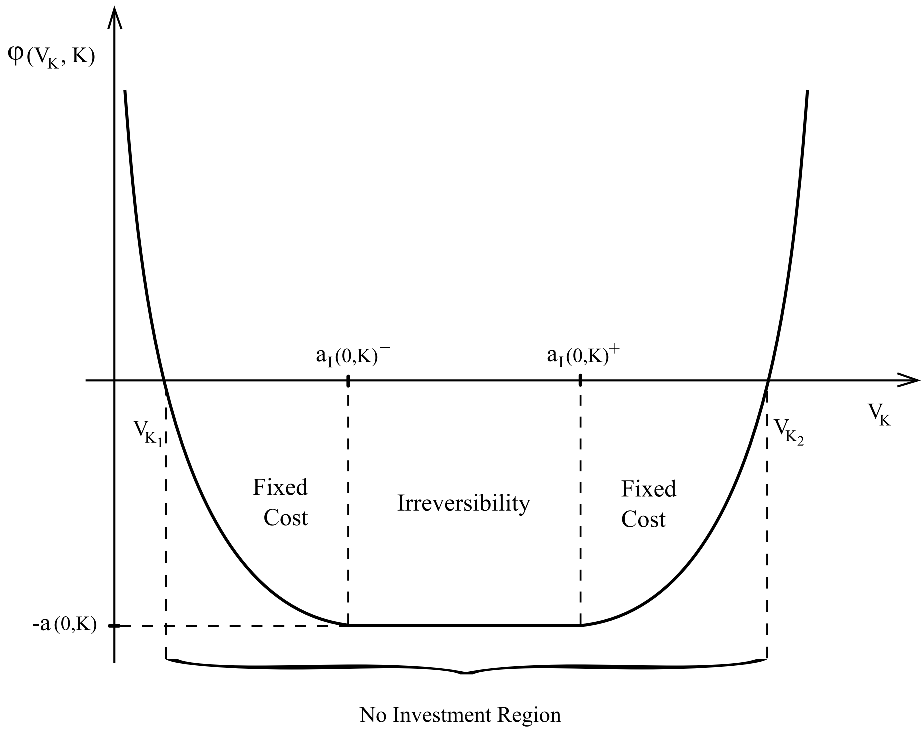

which can store their output may have smaller no-investment regions than firms which cannot store their output, ceteris paribus. To get a better idea of the magnitude of the effect of storage on the investment decision, we solve numerically a discrete-time dynamic optimization problem.

In the next subsection, we describe the discrete-time model. The discrete-time model is similar to the continuous-time model, except that investment is now constrained to be non-negative, and the investment horizon is finite. The proposed methodology used to solve the model numerically then described.

1.3.1. Optimization problem

The discrete-time model consists of a finite horizon production economy with a storable commodity. There is a risk-neutral agent (referred to as "the firm") who is initially endowed with inventory of the commodity, and capital stock used in the production process. In each period, the commodity is produced and the agent decides how much of the commodity to sell and how much to invest in capital stock. The agent is faced with stochastic price or demand shocks.

Our setup involves three state variables: the inventory level (Q), the level of the capital stock (K), and the price of the output (P) or the level of the demand (D). There are two control variables: the quantity sold (S) and the level of investment (I). Before presenting the agent's maximization, we discuss in more detail the two endogenous state variables: the commodity inventory (Q), and the capital stock (K).

1.3.1.1. Commodity inventory. The level of inventory of the commodity carried into the next period (Q_{t+1}) depends on the beginning of period level of inventory (Q_t), on the quantity produced in the current period ($f(K_t)$), and on the quantity sold in the current period (S_t). We assume that that there is no spoilage of inventory. Specifically, the commodity inventory accumulation equation is given by:

$$Q_{t+1} = Q_t + f\left(K_t\right) - S_t.$$

Since the quantity sold in a given period cannot exceed the the quantity on hand during this period, the agent is faced with an inventory constraint:

$$S_t \leq Q_t + f\left(K_t\right).$$

1.3.1.2. Capital stock. The production of the commodity requires the use of capital stock (K_t). The capital stock depreciates over time, but the agent can increase the level of capital stock by investing. The level of the capital stock at the beginning of the next period (K_{t+1}) depends on the level of capital stock in the current period (K_t), the depreciation rate (δ), and the current investment in capital stock (I_t). We assume that there is a one period time-to-build. That is, any investment made in the current period only becomes productive in the following period. The capital accumulation equation is given by:

$$K_{t+1} = K_t(1-\delta) + I_t.$$

In order to increase the level of capital stock, the agent must incur adjustment costs that depends both on the amount invested and on the capital stock in place. We assume that the adjustment cost is of the form:

$$a(I,K) = \begin{cases} I + b\left(\frac{I}{K}\right)^2 K + \Psi K & \text{if} \quad I \geq 0 \\ \infty & \text{if} \quad I < 0. \end{cases}$$

If investment is non-negative, the first two terms of the adjustment cost function represent standard convex adjustment costs, with the first term representing the cost of purchasing capital (with the price of capital normalized to 1). The last term is the fixed cost component of the adjustment cost. The fixed cost is proportional to the capital stock so that the fixed cost component does not become insignificant for large firms. The augmented adjustment cost function is equal to infinity when investment is negative to ensure that investment is always non-negative. This assumption could be relaxed to allow for partial reversibility. While the fixed cost makes investment lumpier, the quadratic cost component smooths investment over time since investment in capital stock reduces future adjustment costs.

1.3.1.3. The agent's maximization problem when price is stochastic. Let T be the terminal date. When the price is stochastic, the state variables are the price of the produced output (P), the inventory level (Q), and the level of the capital stock (K). The control variables are the quantity sold (S) and the level of investment (I). For each time period $t < T$, the firm's maximization problem is:

$$V\left(P_t, Q_t, K_t\right) = \max_{\{S_t, I_t\}} \pi\left(S_t, P_t, Q_t, K_t\right) - a(I_t, K_t)1_{\{I \neq 0\}} + e^{-r} E_t\left[V\left(P_{t+1}, Q_{t+1}, K_{t+1}\right)\right]$$

s.t.

$$Q_{t+1} = Q_t + f\left(K_t\right) - S_t \tag{1.6}$$

$$K_{t+1} = K_t(1-\delta) + I_t \tag{1.7}$$

$$P_{t+1} = P_t + \sigma P_t \varepsilon_{t+1} \tag{1.8}$$

$$S_t \leq Q_t + f\left(K_t\right) \tag{1.9}$$

$$S_t \geq 0 \tag{1.10}$$

$$V(P_T, Q_T, K_T) = h(P_T, Q_T, K_T) \tag{1.11}$$

where $\varepsilon_{t+1} \sim N(0,1)$, σ is the volatility of the price shock, $h(\cdot,\cdot,\cdot)$ is the function representing the terminal condition, $\pi(S_t, P_t, Q_t, K_t)$ is the current period profit function, r is the discount rate, $f(K_t)$ is the production function, and $a(I_t, K_t)$ is the adjustment cost to increase the capital stock.

Equations (1.6), and (1.7) are the accumulation equations for the inventory and capital stock, respectively. Equation (1.6) shows that any quantity of the commodity not sold in the current period is stored and carried over into the next period, and that the inventory does not depreciate. The first inequality constraint (equation (1.9)) represents the fact that the agent cannot sell more of the commodity than he has on hand. We refer to equation (1.9) as the "inventory constraint".

1.3.1.4. The agent's maximization problem when demand is stochastic. Let T be the terminal date. When demand is stochastic, the state variables are the level of the demand (D), the inventory level (Q), and the level of the capital stock (K). The control variables are the quantity sold (S) and the level of investment (I). For each time period $t < T$, the firm's maximization problem is:

$$V(D_t, Q_t, K_t) = \max_{\{S_t, I_t\}} \pi(S_t, D_t, Q_t, K_t) - a(I_t, K_t)1_{\{I \neq 0\}} + e^{-r} E_t \left[V(D_{t+1}, Q_{t+1}, K_{t+1}) \right]$$

s.t.

$$Q_{t+1} = Q_t + f(K_t) - S_t \tag{1.12}$$

$$K_{t+1} = K_t(1-\delta) + I_t \tag{1.13}$$

$$D_{t+1} = D_t + \kappa_D (\phi - D_t) + \sigma_D \varepsilon_{t+1}^D \tag{1.14}$$

$$S_t \leq Q_t + f(K_t) \tag{1.15}$$

$$S_t \geq 0 \tag{1.16}$$

$$V(P_T, Q_T, K_T) = h(P_T, Q_T, K_T) \tag{1.17}$$

where $\varepsilon_{t+1}^D \sim N(0,1)$, κ_D the speed of the mean reversion of the demand, ϕ is the long-run mean of the demand, σ_D is the volatility of the demand shock, $h(\cdot,\cdot,\cdot)$ is the function representing the terminal condition, $\pi(S_t, P_t, Q_t, K_t)$ is the current period profit function, r is the discount rate, $f(K_t)$ is the production function, and $a(I_t, K_t)$ is the adjustment cost to increase the capital stock.

Equations (1.12), and (1.13) are the accumulation equations for the inventory and capital stock, respectively. Equation (1.12) shows that any quantity of the commodity not sold in the current period is stored and carried over into the next period, and that the inventory does not depreciate. The first inequality constraint (equation (1.15)) represents the fact that the agent cannot sell more of the commodity than he has on hand. We refer to equation (1.15) as the "inventory constraint".

1.3.2. Description of numerical procedure used to solve the discrete-time problem

In this subsection, we give a brief overview of the numerical procedure used to solve the discrete time problem. Note that, for expositional purposes, we assume here that the uncertainty in the profit function comes from the output price. The case of stochastic demand is solved similarly.

Quasi-random sequences to generate grid points. In particular, we use the Niederreiter quasi-random sequence. This low discrepancy sequence tends to sample space more uniformly than pseudorandom sequences (Niederreiter (1988)).

The first step is to set up grid points such that all parts of the three-dimensional state space are well represented. To do so, we first find bounds for the state variables. For the output price (P_t), the minimum and maximum values depend on the volatility of the price process. Given the dynamics for the output price, we can shrink the state space for P_t as we move closer to time 0. For the inventory (Q_t), the state space is bounded below by zero and bounded above by the maximal amount which can be sold in the future. Finally, for the capital stock (K_t), the upper bound is given by the amount of capital stock required to produce enough output to meet the demand in every period. We then use a Niederreiter quasirandom sequence to generate a sample of points which is "uniform"over the state space.[1]

The next step is to find the value function at the terminal date T at each grid point

$$V(P_T, Q_T, K_T) = h(P_T, Q_T, K_T).$$

[1]Note that grid points were also generated using Faure and Sobol quasirandom sequences, and using pseudorandom number sequences. This did not affect our results.

We then use neural networks to find an approximate functional form for the value function at time T. (i.e. to find $\hat{V}\left(P_T, Q_T, K_T\right)$). A detailed description of the neural network used is given in Appendix A. Appendix A also provides an overview the steps to follow to ensure that the neural network does not overfit.

Finally, for all the other time periods $0 \leq t < T$, we solve the following maximization problem:

$$V(P_t, Q_t, K_t) = \max_{\{S_t, I_t\}} \pi\left(S_t, P_t, Q_t, K_t\right) - a(I_t, K_t)1_{\{I_t \neq 0\}} + e^{-r} E_t\left[\hat{V}\left(P_{t+1}, Q_{t+1}, K_{t+1}\right)\right]$$

s.t.

$$\begin{aligned}
Q_{t+1} &= Q_t + f\left(K_t\right) - S_t \\
K_{t+1} &= K_t(1-\delta) + I_t \\
P_{t+1} &= P_t + \sigma P_t \varepsilon_{t+1} \\
0 &\leq S_t \leq Q_t + f\left(K_t\right) \\
I_t &\geq 0
\end{aligned}$$

where $\hat{V}\left(P_{t+1}, Q_{t+1}, K_{t+1}\right)$ is the neural net approximation of the value function at time $t+1$ and $E_t\left[V\left(P_{t+1}, Q_{t+1}, K_{t+1}\right)\right]$ is calculated using Gauss-Hermite quadrature.

1.4. Numerical analysis

Before discussing our quantitative results, we will discuss the specific functional forms and parameter values chosen to solve the problem numerically.

The profit function is given by:

$$\pi\left(S, P, Q, K\right) = \beta P \min(S, D)$$

where $\beta \leq 1$, and D is the demand.

The augmented adjustment cost function is defined as:

$$a(I, K) = \begin{cases} I + b\left(\frac{I}{K}\right)^2 K + \Psi K & \text{if} \quad I \geq 0 \\ \infty & \text{if} \quad I < 0. \end{cases}$$

When investment is non-negative, the first two terms represent standard convex adjustment costs, with the first term representing the cost of purchasing capital (with the price of capital normalized to 1). The last term is the fixed cost component of the adjustment cost. The fixed cost is proportional to the capital stock so that the fixed cost component does not become insignificant for large firms. The augmented adjustment cost function is equal to infinity when investment is negative to ensure that investment is always non-negative. This assumption could be relaxed to allow for partial reversibility.

The production function is assumed to depend only on the level of the capital stock in place and to have constant returns to scale. That is,

$$f(K) = K.$$

We assume that, at the terminal date T, the firm sells the quantity on hand of the commodity, up to the demand. As such, the terminal condition is given by:

$$h(P_T, Q_T, K_T) = \beta P_T \min(Q_T + f(K_T), D).$$

To solve the model numerically, we must choose parameter values. Unless otherwise specified, the following parameters will be held constant throughout:

- The discount rate (r) is equal to zero.
- There are no quadratic costs of adjustment (i.e. $b = 0$).
- The initial spot price (S_0) is equal to 10.
- The multiplicative constant in the profit function (β) is set equal to 0.7.
- The long-run level of the demand (ϕ) and the initial level of demand (D_0) are set equal to one.
- When the demand is stochastic, the speed of mean reversion (κ_D) is set equal to 0.5.

1.4.1. Quantitative results - stochastic price

The optimization problem using an adjustment cost function which consists only of the purchasing cost of capital and a fixed cost (i.e. $b = 0$). Since the parameter b in the adjustment cost function is set equal to zero, the threshold value for V_K is an increasing function of K, which means that firms with more capital stock in place will be in the "no investment"region. To compare how the fixed cost affects firms' decisions to invest when there is a fixed adjustment cost, we compare the size of the largest firm to invest when output can be stored to the size of the largest firm to invest when output cannot be stored. Size here is defined as the amount of capital stock in place. Note that whether or not storage is possible, both firms have no inventory at time 0. Also, unless otherwise specified, we assume that the capital stock does not depreciate (i.e. $\delta = 0$).

Table 1.1 presents the results for the two period case, using different levels of the fixed cost parameter (ψ), when the volatility of the price (σ) is 0.10. The numbers in the first two rows represent the size of the largest firm to invest when storage is and is not possible. The third row gives the difference between the size of the largest firm to invest when storage is possible and the size of the largest firm to invest when storage is not possible, divided by the size of the largest firm to invest when storage is not possible. As the fixed cost increases, the investment becomes lumpier since the firm wants to avoid incurring the fixed cost again. The lumpier investment makes bad realizations of the price more costly. As such, the size of the investment region shrinks with the fixed costs. The rate at which the investment region shrinks, however, is lower for firms which can store their output since the effect of the higher fixed cost is dampened by the fact that storage gives firms the option to wait for a high price to sell their output. We refer to this additional flexibility as the "timing option"associated with storage. Since the value of the timing option rises with the fixed cost, the difference between firms with and without storage increases monotonically with the fixed cost.

Table 1.2 shows how the variability of the price shock affects the investment region. Table 1.2 is equivalent to Table 1.1, except that the volatility of the price shock is now lower ($\sigma = 0.01$ instead of 0.10). The comparison of Tables 1.1 and 1.2 shows that the difference in size of the investment region with and without storage is smaller when the volatility of the shock is lower. Furthermore, while the investment region shrinks when the volatility falls when storage is possible, the investment region is not affected by the level of uncertainty when the output cannot be stored. Both of these results can be explained by the fact that reducing the volatility reduces the value of the timing option associated with storage.

Tables 1.3 and 1.4 show the effect of the horizon over the size of the investment region. Tables 1.3 and 1.4 are equivalent to Tables 1.1 and 1.2, respectively, except that $T = 5$ instead of $T = 2$. The results of these two tables show that while the investment region tends to be larger for the longer horizon when storage is possible, the investment region tends to be smaller for the longer horizon when output cannot be stored. The smaller investment region when storage is not possible is due to the fact that, when the horizon is longer, firms want to wait until more uncertainty is resolved before investing. When the output can be stored, however, the value of the timing option associated with stored output rises with the horizon.[2]

Table 1.5 presents the same results as in Table 1.1, except that the capital stock rate depreciates at a rate of 20 percent per period (i.e. $\delta = 0.20$). The numbers in the first two rows represent the size of the largest firm to invest when storage is and is not possible, divided by the maximal amount of capital needed to meet demand in all periods without investing.[3] One noteworthy difference between the results of Tables 1.1 and 1.5 is that, when the capital stock depreciates, the difference between the size of the largest firm to invest when storage is and is not possible is not always monotonically increasing in the fixed cost. Similarly, the gap between the investment region with and without

[2]The intuition for this is the same as that for the value of an American call option rising with the time horizon.

[3]For instance, if a firm cannot store its output, then it will need to have 1.5625 units of capital stock in place at time 0 in order to be able to meet the demand in every period without investing($\frac{1}{0.8^2}$). As such, the first entry in the second row of Table 1.5 (0.8) represents a level of capital stock of 1.25 ($0.8 = \frac{1.25}{1.5625}$), which means that only firms which are not able to meet the demand in the next period will invest. Firms who can meet the demand in the next period will defer their investment since the depreciation lowers the cost of investment in the next period.

depreciation initially narrows as the fixed cost increases, but then starts to increase with the fixed cost when the fixed cost is sufficiently high. This is because depreciation affects the investment decision in two opposite directions. On the one hand, higher depreciation means that more capital is required to meet demand in every period. On the other hand, depreciation makes waiting to invest more attractive since the fixed cost will be lower in the future.[4] For a sufficiently high level of the fixed cost, the benefit of a lower fixed cost which comes from waiting to invest becomes the dominant effect. This causes the difference between the investment regions for firms with and without storage to shrink as the fixed cost increases.

Table 1.1: **Size of largest firm with non-zero investment when price is stochastic, $T = 2$, $\sigma = 0.10$, and $\delta = 0$.** Size is defined as the amount of capital stock in place.

Fixed cost (ψ)	0.0	0.2	0.5	1.0	2.0	3.5	7.0
With storage	1.000	0.976	0.956	0.924	0.856	0.784	0.652
Without storage	1.000	0.974	0.948	0.916	0.820	0.692	0.524
Percent diff.	0.00	0.21	0.84	0.87	4.39	13.29	24.43

Table 1.2: **Size of largest firm with non-zero investment when price is stochastic, $T = 2$, $\sigma = 0.01$, and $\delta = 0$.** Size is defined as the amount of capital stock in place.

Fixed cost (ψ)	0.0	0.2	0.5	1.0	2.0	3.5	7.0
With storage	1.000	0.976	0.952	0.922	0.854	0.7312	0.582
Without storage	1.000	0.974	0.948	0.916	0.820	0.692	0.524
Percent diff.	0.00	0.21	0.40	0.66	4.25	5.66	11.07

[4]Recall that the fixed cost is proportional to the capital stock in place.

Table 1.3: **Size of largest firm with non-zero investment when price is stochastic, $T = 5$, $\sigma = 0.10$, and $\delta = 0$.** Size is defined as the amount of capital stock in place.

Fixed cost (ψ)	0.0	0.2	0.5	1.0	2.0	3.5	7.0
With storage	1.000	0.980	0.960	0.928	0.880	0.864	0.828
Without storage	1.000	0.964	0.940	0.820	0.782	0.692	0.532
Percent diff.	0.00	1.66	2.13	13.17	12.53	24.86	55.64

Table 1.4: **Size of largest firm with non-zero investment when price is stochastic, $T = 5$, $\sigma = 0.01$, and $\delta = 0$.** Size is defined as the amount of capital stock in place.

Fixed cost (ψ)	0.0	0.2	0.5	1.0	2.0	3.5	7.0
With storage	1.000	0.976	0.957	0.884	0.845	0.749	0.579
Without storage	1.000	0.964	0.941	0.819	0.782	0.692	0.534
Percent diff.	0.00	1.23	1.72	7.92	8.03	8.24	8.42

Table 1.5: **Size of largest firm with non-zero investment when price is stochastic, $T = 2$, $\sigma = 0.10$, and $\delta = 0.2$.** Size is defined as the amount of capital stock in place.

Fixed cost (ψ)	0.0	0.2	0.5	1.0	2.0	3.5	7.0
With storage	1.000	0.892	0.888	0.880	0.836	0.764	0.615
Without storage	0.800	0.772	0.740	0.691	0.656	0.391	0.368
Percent diff.	25.00	15.56	20.01	27.30	27.46	95.36	67.21

1.4.2. Quantitative results - stochastic demand

We now assume that the price is constant over time, but that the demand is stochastic.

The only difference with the previous case is that:

$$D_{t+1} = D_t + \kappa_D \left(\phi - D_t\right) + \sigma_D \varepsilon_{t+1}^D$$

where $\varepsilon_{t+1}^D \sim N(0,1)$.

In Tables 1.6 to 1.9 show the size of the largest firm to invest using different parameter values. To analyze the case when demand is stochastic, we assume that the capital stock depreciates at a rate of 20 percent per period (i.e. $\delta = 0.2$).

Tables 1.6 and 1.8 report the size of the largest firm to invest when storage is and is not possible, using different levels of the fixed cost parameter (ψ) and the volatility of the demand (σ_D), when the horizon is two periods. From these tables, we see that, for a given level of the fixed cost, the effect of introducing more volatility on the size of the investment region depends on whether or not the output can be stored.

Table 1.6 shows that if the output can be stored, then introducing more volatility initially shrinks the investment region. The effect of introducing more uncertainty, however, reverses as the volatility continues to increase. The volatility level at which the reversal occurs depends on the level of the fixed cost. The higher the fixed cost, the higher the level of the volatility at which the reversal occurs. Table 1.7 shows that, when output cannot be stored, an increase in uncertainty has the opposite effect. Introducing more volatility initially widens the investment region. The effect of introducing more uncertainty, however, reverses as the volatility continues to shrink. The volatility level at which the reversal occurs depends on the level of the fixed cost. But contrary to the case when output can be stored, the lower the fixed cost, the higher the volatility level at which the reversal occurs.

These results can be understood using the “bad news principle”. When considering whether or not to invest, the firm considers the costs involved under each alternative. If the firm decides to invest, then it must absorb any losses associated with a future fall in demand. If the firm decides not to invest, then it must absorb the cost of unmet demand associated with a future rise in demand. When storage is possible, the costs associated with a fall in demand are greater than the costs of unmet demand, due to the irreversibility and the lumpiness of investment[5], when volatility is low. The costs of a fall in demand, however, are mitigated by the fact that any output produced not sold in the current period can be stored and carried over into the next period. As such, a rise in volatility eventually raises the costs of unmet demand by more than the costs of a fall in demand. When output cannot be stored, the costs of unmet demand are greater than the costs of a fall in demand when volatility is low, due to the absence of buffer stocks and the one-period time to built.[6] The costs of a fall in demand, however, eventually increase at a faster rate than the cost of unmet demand due to the lumpiness and irreversibility of investment when volatility increases.

The result for the case when output can be stored can also be understood in the context of (real) option theory. As we introduce more uncertainty, the option to wait becomes more valuable. The flexibility which comes from holding inventory, however, also becomes more valuable as uncertainty increases. If we think of the value placed on this flexibility as a dividend which accrues to the holder of the inventory, the result from Table 1.6 is analogous to the early exercise of an American call option when dividends are sufficiently high.[7] In this case, the call option is the option to invest, and the strike price is the cost of investment.

Tables 1.8 and 1.9 present the same results as in Tables 1.6 and 1.7, except that the horizon is now five periods instead of two. Qualitatively similar results hold for the

[5]Recall that the fixed cost makes investment lumpier because the firm wants to minimize the number of times that it invests

[6]Recall that there is a one-period lag between investment and the productive use of the new capital stock.

[7]The dividend which accrues to the holder of inventory is often referred to as the "convenience yield".

Table 1.6: **Size of largest firm with non-zero investment when demand is stochastic, $T = 2$ and output can be stored.** Size is defined as the amount of capital stock in place.

	volatility (σ)						
Fixed cost (ψ)	0.001	0.05	0.10	0.15	0.20	0.25	0.30
0.00	1.3888	1.2531	1.2836	1.3112	1.3701	1.3946	1.4085
0.10	1.2247	1.1736	1.2069	1.2153	1.2440	1.2542	1.2711
0.50	1.1900	1.1448	1.1442	1.1354	1.1539	1.1639	1.1696
0.80	1.1702	1.1240	1.1153	1.1145	1.1087	1.1038	1.1100
1.00	1.1405	1.1040	1.0954	1.0952	1.0819	1.0737	1.0833
2.00	1.0909	1.0664	1.0435	1.0236	1.0128	0.9935	0.9920

longer time horizon. The main difference is that the “reversal”occurs at a lower level of volatility in the five period case. When the output can be stored, the reversal occurs more quickly because the longer horizon increases the value of the flexibility which comes from holding inventory. As such, the option to invest gets exercised earlier. When storage is not possible, the reversal occurs more quickly when the horizon is longer because the cost of a fall in demand is higher due to lumpiness of investment and the depreciation of the capital stock.

1.5. Conclusion

In this chapter, we examined the effect of including inventory as a state variable in capital investment models with non-convex adjustment costs. Our analysis focuses on understanding the determinants of the investment region. We find that the possibility of storage can significantly affect the investment behavior of firms, both quantitatively and qualitatively. Our results are obtained assuming that there are no storage costs. Adding a storage cost would decrease the impact of storage on the investment decision.

Table 1.7: **Size of largest firm with non-zero investment when demand is stochastic, $T = 2$ and output cannot be stored.** Size is defined as the amount of capital stock in place.

	volatility (σ)						
Fixed cost (ψ)	0.001	0.05	0.10	0.15	0.20	0.25	0.30
0.00	1.2504	1.3640	1.4997	1.5668	1.5683	1.5691	1.5692
0.10	1.2247	1.2450	1.3045	1.3511	1.4151	1.4447	1.4698
0.50	1.1731	1.1795	1.1999	1.2073	1.2170	1.2181	1.2148
0.80	1.1652	1.1676	1.1720	1.1754	1.1719	1.1640	1.1496
1.00	1.1646	1.1657	1.1581	1.1574	1.1539	1.1339	1.1164
2.00	1.0959	1.0942	1.0745	1.0635	1.0458	1.0236	0.9950

Table 1.8: **Size of largest firm with non-zero investment when demand is stochastic, $T = 5$ and output can be stored.** Size is defined as the amount of capital stock in place.

	volatility (σ)						
Fixed cost (ψ)	0.001	0.05	0.10	0.15	0.20	0.25	0.30
0.00	1.8592	1.2979	1.4488	1.5489	1.6454	1.6793	1.7058
0.10	1.2381	1.2294	1.2406	1.3066	1.3355	1.4334	1.5496
0.50	1.1960	1.1842	1.1729	1.1765	1.1768	1.2344	1.2428
0.80	1.1899	1.1722	1.1435	1.1405	1.1528	1.1598	1.2182
1.00	1.1817	1.1598	1.1176	1.1118	1.1107	1.1151	1.1498
3.50	1.1568	1.0492	0.9877	0.9793	0.9781	0.9692	0.9745

If the storage cost is sufficiently large, then there is no difference between the investment decision of firms that can store their output and firm that cannot.

Our results represent a first step towards understanding the investment behavior of firms producing storable commodities. Understanding the investment behavior of such firms is important to gain more insight into the dynamics of spot and futures equilibrium prices of storable commodities. Such an equilibrium model is proposed in the next chapter.

Table 1.9: **Size of largest firm with non-zero investment when demand is stochastic, $T = 5$ and output can not be stored.** Size is defined as the amount of capital stock in place.

	volatility (σ)						
Fixed cost (ψ)	0.001	0.05	0.10	0.15	0.20	0.25	0.30
0.00	1.2502	1.3633	1.5041	1.6391	1.7407	1.8710	1.9733
0.10	1.2352	1.2530	1.3105	1.3855	1.4547	1.5129	1.6081
0.50	1.2060	1.2194	1.2199	1.2225	1.2245	1.2279	1.2574
0.80	1.2025	1.2058	1.1445	1.1401	1.1253	1.1118	1.1113
1.00	1.1438	1.1822	1.0984	1.0367	1.0257	1.0222	1.0214
3.50	0.9544	0.7884	0.6557	0.5611	0.5015	0.4652	0.4247

CHAPTER 2

Prices of Storable Commodities with Irreversible Investment and Liquidity Constraints

2.1. Introduction

Over the past decade, the number of investable commodity indices and commodity-linked investment instruments has grown significantly, as has the trading volume for these securities. For instance, as Figure 2.1 shows, the average daily volume for options and futures traded on COMEX and NYMEX has almost been rising monotonically since 1994, almost quadrupling in the last 13 years. As such, interest in models which lay down the foundations for the dynamics of commodity markets has risen. Also, due to an increase in the variety of underlying commodities, it has become increasingly important to understand how the drivers of the price process differ across types of commodities.

The purpose of this chapter is to model the price dynamics of a commodity in a production economy framework. Our model can be used to study the price dynamics not only of industrial commodities (such as copper and aluminum), but also of commodities which serve both as industrial commodities and stores of value (such as gold and silver). The model is an equilibrium storage model which incorporates the production side of the economy. With production, the agent can respond to shocks not only through storage, but also through investment in capital stock, which depreciates over time. While the depreciation of the capital stock increases the role of storage, the investment in capital stock makes stockouts more likely in periods of high investment. In this framework, the spot price dynamics are determined by both the storage decision and the demand and production shock processes.

Our model differs significantly from existing commodity pricing models because it allows the primary purpose served by a commodity to vary over time. While it is not a new result that the price of commodities which serve primarily as stores of value behaves differently than the price of industrial commodities (for instance, Brennan (1991) and Casassus and Collin-Dufresne (2005)), existing models do not seem to adequately capture the price dynamics of a commodity, such as gold, for which the importance of the industrial role relative to the store of value role varies over time (for instance, Brennan (1990) and Schwartz (1997)). The focus of our analysis is on how the nature of a commodity affects the joint dynamics of the spot price and convenience yield.[1] As such, we contribute to the literature by providing a theoretical framework to understand the determinants of the convenience yield for different types of commodities over various horizons. This is important since we would expect the investment decision of commodity producing firms to be driven, at least partially, by the convenience yield of the commodity over the investment horizon. Finally, although not directly addressed in this paper, our work represents an important first step to understanding the benefits of including different types of commodities in an investor's portfolio.

[1]The convenience yield can be thought of as the benefit from physical ownership of the commodity (as opposed to a financial claim to the commodity). More technically, it is the difference between the risk-free rate and the return on the commodity inferred from the forward price. In the spirit of the original theory of storage (Kaldor (1939), Working (1949), Telser (1958), Brennan (1958)), our measure of convenience captures the timing option which comes from immediate ownership of the commodity.

To study the behavior of commodity prices, researchers in the commodity pricing literature have used either a statistical reduced-form approach or an equilibrium approach. While the reduced-form approach is useful to identify the dynamics which seem to best capture the features of a given commodity market, such an approach is limited in that it does not make any predictions as to how a change in a parameter estimate affects other parameters (for instance, Gibson and Schwartz (1990), Schwartz (1997), Schwartz and Smith (2000) and Casassus and Collin-Dufresne (2005)). To gain a better understanding of the drivers of commodity prices, competitive equilibrium storage models are used (for instance, Scheinkman and Schechtman (1983), Deaton and Laroque (1992), Deaton and Laroque (1996), Chambers and Bailey (1996) and Routledge, Seppi, and Spatt (2000)).[2] Because the focus of these models is on the actions of risk-neutral profit maximizing inventory holders rather than on the actions of producers, such models are not intended to give storage a significant role in the determination of longer term prices. As a result, the price dynamics which come out of these models are determined primarily by the chosen processes for the production and demand shocks.

The dynamics in existing competitive storage models is such that periods of scarcity of the commodity (i.e. periods of low inventory) correspond to periods of high convenience yield and high spot price. As such, these models predict a positive relation between the spot price and the convenience yield. This result is intuitive and seems to hold for industrial commodities such as copper, aluminum and crude oil (for instance, see Figure

[2]More recently, researchers have also used commodity pricing equilibrium models in a production economy framework (for instance, Carlson, Khokher, and Titman (2002), Casassus, Collin-Dufresne, and Routledge (2004) and Kogan, Livdan, and Yaron (2005)). Such models, however, do not allow the final output to be stored and, as such, are not as directly comparable to our model as are competitive equilibrium storage models.

2.2). This relation, however, seems to break down for gold and, to a lesser extent, silver. Figure 2.3 shows the relation between the spot price and the one-year convenience yield of gold from 1984 to 2005. While there do appear to be periods during which the spot price and the convenience yield of gold move in the same direction, there are other periods during which the spot price and the convenience yield clearly move in opposite directions. The non-monotonic relation between the spot price and the convenience yield for gold and silver is precisely the feature of the data that our model sets out to explain.

The model is an infinite horizon production economy with a risk-averse agent and two goods: a liquid asset and a storable commodity. The agent gets utility by consuming from liquid wealth. In each period, the agent produces the commodity. He then decides how much of the commodity to sell in exchange for liquid wealth, how much liquid wealth to invest in capital stock, and how much liquid wealth to consume. The agent is faced production shocks, demand shocks, and consumption shocks. In this framework, there will be a positive relation between the spot price and the convenience yield as long as the agent has enough liquid wealth to make his storage decision independently from his short-term consumption decision. When liquidity is tight, the storage decision is no longer independent from the consumption decision since, in this case, the agent considers his short-term consumption needs when deciding how much of the commodity to sell. In this case, there can be a negative relation between the spot price and the convenience yield since the demand for liquidity puts downward pressure on the price and upward pressure on the convenience yield.[3]

[3]The downward pressure on the price is due to the fact that, ceteris paribus, the agent sells more of the commodity when he is liquidity constrained due to the liquidity provided by sale of the commodity. The upward pressure on the convenience yield is due to the increased value placed on the timing option which comes from physical ownership of the commodity.

The model is solved numerically using the value function iteration procedure. At each step, the value function is approximated using feedforward neural networks. Because neural networks are able to capture the underlying systematic aspects of the data, they are powerful functional approximators. Also, while most of the more commonly used approximation schemes (for instance, polynomials and splines) are efficient for lower dimensional problems, these schemes are subject to the curse of dimensionality, with the computational requirements often growing geometrically with the dimensionality of the problem. Using neural networks, our methodology can straightforwardly be adjusted to allow for additional dimensions.

The chapter is organized as follows. In Section 2, we discuss our model setup. In Section 3, we present our numerical analysis. In Section 4, we discuss the empirical implications of our model. In Section 5, we empirically test these implications. Finally, in Section 6, we conclude and discuss paths for future research.

2.2. Description of the model

In this section, we describe our model in general terms. Our setup allows us to examine the joint dynamics of the spot price and convenience yield when the primary purpose served by a commodity varies over time.

2.2.1. Model setup

Our model consists of an infinite horizon economy with two goods: a liquid asset and a commodity. There is a risk-averse agent, who is initially endowed with inventory of the commodity, a stock of the liquid asset (his liquid wealth), and capital stock used to

Figure 2.1: Average daily volume for futures and option contracts traded on NYMEX and COMEX. This figure shows the average daily volume for futures and option contracts traded on NYMEX and COMEX. The data was obtained from the NYMEX web site.

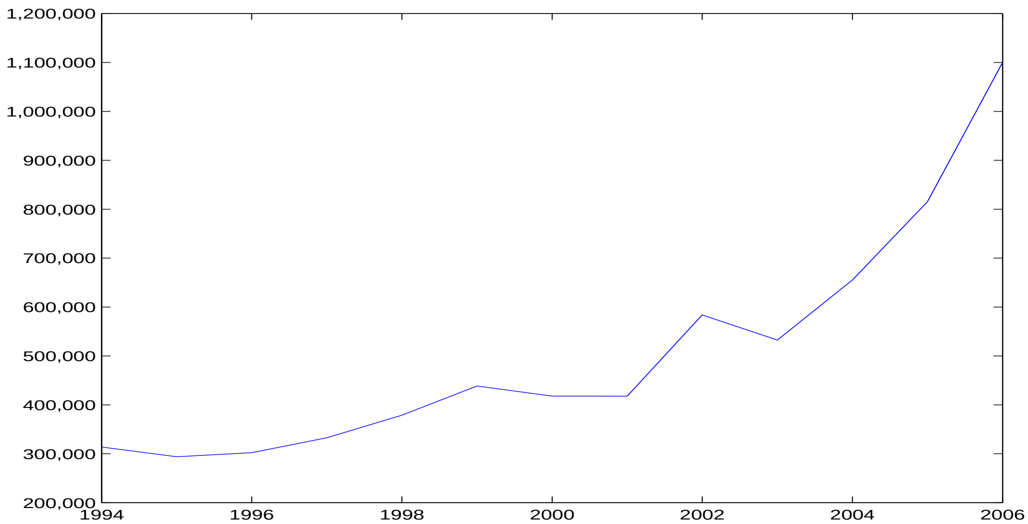

produce the commodity. The agent gets utility only by consuming from his liquid wealth over time. In each period, the commodity is produced and the agent decides how much of the commodity on hand to sell in exchange for liquid wealth[4], how much of his liquid wealth to consume, and how much of his liquid wealth to invest in capital stock. The agent is faced with production, demand, and consumption shocks.

Our setup involves six state variables: the level of inventory of the commodity (Q_t), the level of liquid wealth (W_t), the level of capital stock (K_t), the production shock (ε_t), the demand shock (z_t), and the consumption shock (ς_t). There are three control variables: the commodity sold (S_t), the investment in capital stock (I_t), and the consumption of liquid wealth (C_t). Before presenting the agent's maximization, we discuss in more detail

[4]Technically, we are assuming that there exists a conversion technology which transforms the commodity into liquid wealth. However, the "selling", as opposed to "converting", terminology will be used throughout because of its intuitive appeal.

Figure 2.2: Spot price and the one-year convenience yield of copper. This figure shows the spot price of copper (defined as the nearest-to-maturity futures price) and the one-year convenience yield obtained from the one-year copper futures contract using the one-year constant maturity Treasury rate from 1990 to 2005. Data for futures prices and Treasury rates are obtained from Datastream. Copper futures contracts used are those traded on NYMEX. The one-year convenience yield at time t is computed as $q_{t,t+1} = \left(\frac{F_{t,t+\varepsilon}}{F_{t,t+\varepsilon+1}}(1+r_1) - 1\right)$ where $F_{t,t+\varepsilon}$ is the nearest-to-maturity futures contract and r_1 is the one-year constant-maturity Treasury rate.

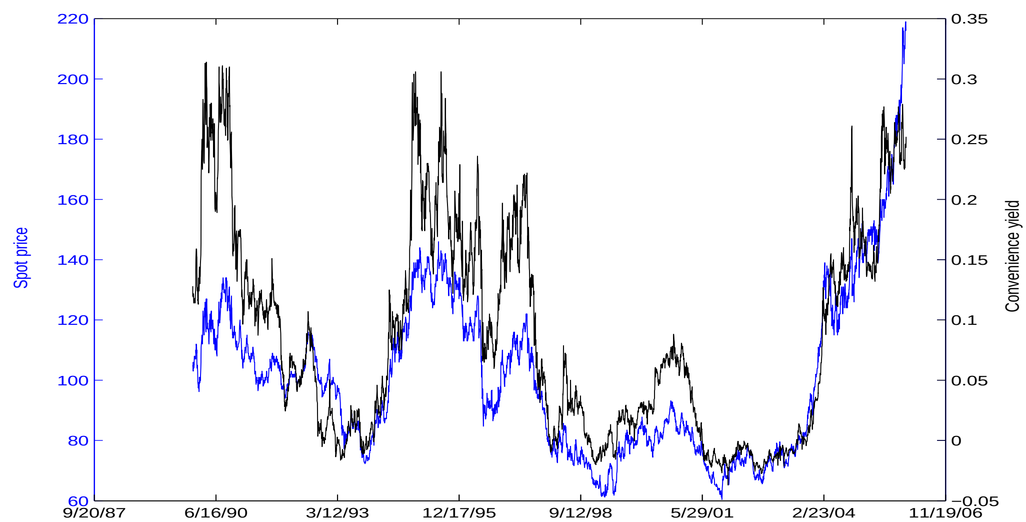

the three endogenous state variables: the commodity inventory (Q_t), the capital stock (K_t), and the liquid wealth (W_t).

2.2.1.1. Commodity inventory. The level of inventory of the commodity carried into the next period (Q_{t+1}) depends on the beginning of period level of inventory (Q_t), the quantity produced in the current period ($f(\varepsilon_t, K_t)$), and the quantity sold (S_t). We assume that that there is no cost of storage nor spoilage of inventory. The commodity inventory accumulation equation is given by:

$$Q_{t+1} = Q_t + f(\varepsilon_t, K_t) - S_t.$$

Figure 2.3: Spot price and the one-year convenience yield of gold. This figure shows the spot price of gold (defined as the nearest-to-maturity futures price) and the one-year convenience yield obtained from the one-year gold futures contract using the one-year constant maturity Treasury rate from 1984 to 2005. Data for futures prices and Treasury rates are obtained from Datastream. Gold futures contracts used are those traded on NYMEX. The one-year convenience yield at time t is computed as $q_{t,t+1} = \left(\frac{F_{t,t+\varepsilon}}{F_{t,t+\varepsilon+1}}(1+r_1)-1\right)$ where $F_{t,t+\varepsilon}$ is the nearest-to-maturity futures contract and r_1 is the one-year constant-maturity Treasury rate.

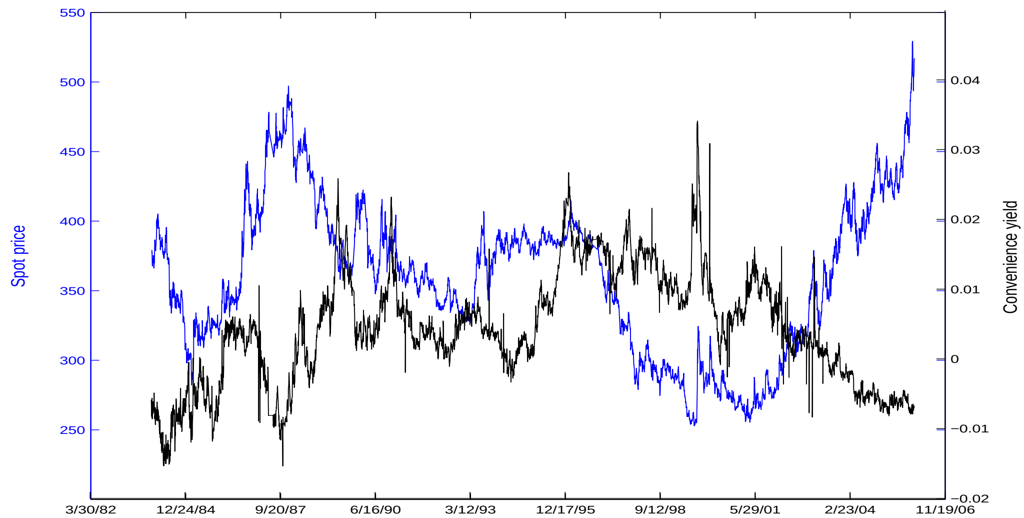

As shown by this equation, any inventory of the commodity not sold in the current period is carried over to the next period.

In any period, the agent cannot sell more of the commodity than the quantity on hand. As such, the is faced with the following inventory constraint:

$$S_t \leq Q_t + f(\varepsilon_t, K_t).$$

2.2.1.2. Capital stock. It is important to incorporate the production side of the economy because different aspects of the production process result in different commodity price dynamics. For instance, the cost of investment in capital stock affects the importance of storage in the determination of longer term prices.

The production of the commodity requires the use of capital stock (K_t). The capital stock depreciates over time, but the agent can increase the level of capital stock by investing. The level of the capital stock at the beginning of the next period (K_{t+1}) depends on the level of capital stock in the current period (K_t), the depreciation rate (δ), and the current investment in capital stock (I_t). We assume that there is a one period time-to-build. That is, any investment made in the current period only becomes productive in the following period. The capital accumulation equation is given by:

$$K_{t+1} = K_t(1 - \delta) + I_t.$$

In order to increase the level of capital stock, the agent must incur adjustment costs that depend both on the amount invested and on the capital stock in place. More specifically, the adjustment cost is of the form:

$$a(I_t, K_t) = \begin{cases} cI_t + b\left(\frac{I_t}{K_t}\right)^2 K_t & \text{if } I_t > 0 \\ 0 & \text{if } I_t = 0 \\ \infty & \text{if } I_t < 0. \end{cases}$$

This form of the adjustment cost function means that investment is irreversible, and that the agent must incur linear and quadratic costs to increase the level of the capital stock. The quadratic cost component smooths investment over time since investment in capital stock reduces future adjustment costs.

2.2.1.3. Liquid wealth. In each period, the agent consumes from liquid wealth. The level of liquid wealth at the beginning of the next period (W_{t+1}) depends on the level of the liquid wealth at the beginning of the current period (W_t), the proceeds from the sale of the commodity in the current period ($p_t S_t$), the adjustment costs incurred to invest in capital stock ($a(I_t, K_t)$), and the quantity consumed (C_t). The liquid wealth accumulation equation is given by:

$$W_{t+1} = W_t + p_t S_t - a\left(I_t, K_t\right) - C_t.$$

As shown by this equation, any liquid wealth not consumed in the current period is carried over to the next period.

An important assumption that we make is that the agent is faced with a liquidity constraint. Specifically, the agent cannot consume more than the amount currently on hand:

$$C_t \leq W_t + p_t S_t - a\left(I_t, K_t\right).$$

This liquidity constraint can be interpreted as a borrowing constraint. It represents the agent's inability to borrow to increase his current level of consumption. As long as the liquidity constraint is not binding, the liquid wealth account functions as a money market account and accumulates at a rate of zero. When the liquidity constraint is binding, however, the risk-free rate is positive.

The liquidity constraint is economically significant because it affects the determinants of the price dynamics. When the agent has a sufficiently large amount of liquid wealth on hand, the commodity storage decision is made independently from the consumption decision. In this case, the storage decision depends only on the demand and production shocks, and on the level of the commodity inventory. When liquidity is tight, however, the agent also considers his short term demand for liquidity when deciding how much of the commodity to sell. As such, the storage and consumption decisions are no longer independent. As will be explained in greater detail in the following section, it is this lack of independence between the storage and consumption decisions which can give rise to a negative relation between spot price and convenience yield.

2.2.1.4. Agent's maximization problem. Combining all of these constraints, we have the following maximization problem for the agent:

$$V(\Lambda_t) = \max_{S_t, I_t, C_t} u(C_t, \varsigma_t) + \beta E_t \left[V(\Lambda_{t+1}) \right] \tag{2.1}$$

s.t.

$$Q_{t+1} = Q_t + f(\varepsilon_t, K_t) - S_t \tag{2.2}$$

$$K_{t+1} = K_t(1-\delta) + I_t \tag{2.3}$$

$$W_{t+1} = W_t + p_t S_t - a(I_t, K_t) - C_t \tag{2.4}$$

$$S_t \leq Q_t + f(\varepsilon_t, K_t) \tag{2.5}$$

$$C_t \leq W_t + p_t S_t - a(I_t, K_t) \tag{2.6}$$

where $\Lambda_t = (z_t, \varepsilon_t, \varsigma_t, Q_t, K_t, W_t)$, $u(C_t, \varsigma_t)$ is the agent's utility function, β is the subjective discount rate, $p_t S_t$ is the proceeds from the sale of the commodity in period t (i.e. the unit price multiplied by the quantity sold), $f(\varepsilon_t, K_t)$ is the production function, and $a(I_t, K_t)$ is the adjustment cost to increase the capital stock.

Equations (2.2), (2.3) and (2.4) are the accumulation equations for the inventory of the commodity, capital stock and liquid wealth, respectively. The first inequality constraint (equation (2.5)) represents the fact that the agent cannot sell more of the commodity than he has on hand. The second inequality constraint (equation (2.6)) represents the fact that the agent cannot consume more than the liquid wealth on hand. That is, he cannot borrow to increase his consumption in the current period. As will be discussed later, when this constraint is binding, the rate at which the agent sells the commodity depends, at least partially, on the agent's short term consumption needs. This can give rise to a negative relation between the spot price and the convenience yield of the commodity. We refer to equation (2.5) as the "inventory constraint", and to equation (2.6) as the "liquidity constraint".

The above maximization problem can be re-written

$$V\left(\Lambda_t\right) = \max_{S_t, I_t, C_t} u\left(C_t, \varsigma_t\right) + \beta E_t\left[V\left(\Lambda_{t+1}\right)\right] - \lambda_t\left(S_t - Q_t - f\left(K_t, \varepsilon_t\right)\right) - \mu_t\left(C_t - W_t - p_t S_t\right)$$

s.t.

$$\begin{aligned} Q_{t+1} &= Q_t + f\left(\varepsilon_t, K_t\right) - S_t \\ K_{t+1} &= K_t(1-\delta) + I_t \\ W_{t+1} &= W_t + p_t S_t - a\left(I_t, K_t\right) - C_t \end{aligned}$$

where λ_t and μ_t are the Lagrangian multipliers for the inventory and the liquidity constraints, respectively. The first order conditions for the commodity sold (S_t), the consumption (C_t), and the investment in capital stock (I_t) are

$$p_t\left(\beta E_t\left[V_W\left(\Lambda_{t+1}\right)\right]+\mu_t\right) = \beta E_t\left[V_Q\left(\Lambda_{t+1}\right)\right]+\lambda_t \tag{2.7}$$

$$u'\left(C_t\right) = \beta E_t\left[V_W\left(\Lambda_{t+1}\right)\right]+\mu_t \tag{2.8}$$

$$\beta E_t\left[V_K\left(\Lambda_{t+1}\right)\right] = \beta E_t\left[V_W\left(\Lambda_{t+1}\right)\right]\left(c+2b\left(\frac{I_t}{K_t}\right)\right). \tag{2.9}$$

We also have

$$V_Q\left(\Lambda_t\right) = \beta E_t\left[V_Q\left(\Lambda_{t+1}\right)\right]+\lambda_t \tag{2.10}$$

$$V_W\left(\Lambda_t\right) = \beta E_t\left[V_W\left(\Lambda_{t+1}\right)\right]+\mu_t. \tag{2.11}$$

Combining equations (2.7), (2.8) and (2.10), and solving forward, we have

$$p_t = E_t\left[\beta^n\frac{u'\left(C_{t+n}\right)}{u'\left(C_t\right)}p_{t+n}\right]+\sum_{i=0}^{n-1}\beta^i\frac{E_t\left[\lambda_{t+i}\right]}{u'\left(C_t\right)}. \tag{2.12}$$

Combining equations (2.8) and (2.11) and solving forward, we have

$$u'(C_t) = \beta^{n+1}E_t\left[V_W\left(\Lambda_{t+n+1}\right)\right]+\sum_{i=0}^{n}\beta^iE_t\left[\mu_{t+i}\right]. \tag{2.13}$$

Equation (2.12) shows that the spot price depends on the expected marginal benefit from an additional unit of commodity inventory in future periods. Equation (2.13) shows that the consumption decision depends on the expected marginal benefit from an additional unit of liquid wealth in future periods. Equation (2.12) will be useful to derive an expression for the convenience yield and equation (2.13) will be useful to derive and expression for the risk-free rate. More details are given in Appendix B.

2.2.1.5. A note on portfolio choice. In each period, the agent's optimization problem consists of determining the amount of liquid wealth to consume and the best way to move the remaining total wealth through time. In this economy, there are three means by which to move wealth through time. The agent can move wealth by storing the commodity, by investing in capital stock, and by using the liquid wealth account. As such, the agent's maximization problem could be reformulated as a portfolio choice problem where, in each period, he decides simultaneously how much to consume and how the remaining wealth should be allocated between three risky assets (i.e. how much of the commodity to store, how much of the liquid wealth to invest in capital stock, and how much liquid wealth to carry over to the next period). The riskiness of each means of wealth transfer depends primarily on the stochastic nature of the demand, production and consumption shocks.

2.2.1.6. Risk-free rate. In this economy, the risk-free rate depends on the agent's marginal rate of substitution. As long as the liquidity constraint is not binding, the risk-free rate is given by the accumulation rate of the liquid wealth, which is assumed to be zero (see equation (2.4)). When the liquidity constraint is binding, the agent would like to increase his immediate consumption, but is unable to do so. The rate required to make the agent indifferent between consuming an additional unit today and consuming an additional unit in the future is thus higher when there is a liquidity stockout.

This can be seen by writing the n-period risk-free rate, $r_{t,t+n}$, as a function of the Lagrangian multipliers for the liquidity constraint from t up to $t+n$. Using equation

(2.13), we have

$$\begin{aligned}\frac{1}{\left(1+r_{t,t+n}\right)^n} &= E_t\left[\beta^n \frac{u'\left(C_{t+n}\right)}{u'\left(C_t\right)}\right] \\ &= E_t\left[\frac{\beta^n\left(\beta E_{t+n}\left[V_W\left(\Lambda_{t+n+1}\right)\right]+E_{t+n}\left[\mu_{t+n}\right]\right)}{\beta^{n+1}E_t\left[V_W\left(\Lambda_{t+n+1}\right)\right]+\sum_{i=0}^{n}\beta^i E_t\left[\mu_{t+i}\right]}\right] \\ \Rightarrow \quad r_{t,t+n} &= \left(1+\frac{\sum_{i=0}^{n-1}\beta^i E_t\left[\mu_{t+i}\right]}{\beta^n E_t\left[u'\left(C_{t+n}\right)\right]}\right)^{\frac{1}{n}} - 1.\end{aligned}$$

Since μ_t represents the shadow value of consuming an additional unit of liquid wealth at time t, this expression shows that the n-period risk-free rate is monotonically increasing in the value associated with liquidity from t to $t+n$.

2.2.1.7. Prepaid forward price and convenience yield. The price, $F^P_{t,t+n}$, at time t, of a note which entitles the holder to a unit of the commodity in n periods is given by

$$F^P_{t,t+n} = E_t\left[\beta^n \frac{u'\left(C_{t+n}\right)}{u'\left(C_t\right)} p_{t+n}\right], \tag{2.14}$$

where p_{t+n} is the commodity spot price at time $t+n$. We refer to $F^P_{t,t+n}$ as the n-period shadow "prepaid forward price". To understand where the expression for $F^P_{t,t+n}$ comes from, consider equation (2.12). This equation shows that the commodity spot price is calculated as the discounted value of the commodity in n periods, plus the marginal benefit of holding the commodity from t to $t+n-1$. Since the holder of an n-period prepaid forward contract only obtains the commodity at $t+n$, the prepaid forward price reflects only the discounted value of the commodity in n periods. The n-period shadow forward price, $F_{t,t+n}$, is obtained by accumulating $F^P_{t,t+n}$ using the n-period risk-free rate:

(2.15)
$$F_{t,t+n} = F^P_{t,t+n} \left(1 + r_{t,t+n}\right)^n .$$ [5]

In each period, the agent cannot sell more of the commodity than he has on hand. That is, he cannot short sell his inventory. This gives rise to a convenience yield, a benefit associated with the physical ownership of the commodity. The convenience yield captures the timing option which comes from the immediate ownership of the commodity and depends on the identity of the holder. In our model, this represents the agent's ability to sell the commodity when it is optimal for him to do so. If commodity stockouts did not affect the agent's ability to sell the commodity[6], then the convenience yield would be zero and the spot price would always be equal to the future asset value of the commodity (i.e. the spot price would be equal to the prepaid forward price).

The relation between the n-period prepaid forward price and the n-period convenience yield, $q_{t,t+n}$, can be seen by combining equations (2.12) and (2.14):

$$\begin{aligned} F^P_{t,t+n} &= \beta^n E_t \left[\frac{u'(C_{t+n})}{u'(C_t)} p_{t+n} \right] \\ &= p_t - \sum_{i=0}^{n-1} \beta^i \frac{E_t\left[\lambda_{t+i}\right]}{u'(C_t)} \\ &= \frac{p_t}{\left(1 + q_{t,t+n}\right)^n} \end{aligned}$$

[5] The n-period forward price is the price determined at time t for one unit of the commodity at time $t+n$, to be paid at time $t+n$.

[6] For instance, if the agent could short sell the commodity.

where $q_{t,t+n} = \left(\frac{p_t u'(C_t)}{p_t u'(C_t) - \sum_{i=0}^{n-1} \beta^i E_t[\lambda_{t+i}]}\right)^{\frac{1}{n}} - 1$. The expression for $q_{t,t+n}$ shows that the n-period convenience yield depends on both the likelihood and the severity of a commodity stockout prior to the maturity of the forward contract.

2.3. Numerical analysis

In this section, we solve the model numerically to study the price dynamics of a commodity which can serve both as an industrial commodity and a store of value. The numerical procedure used is the value function iteration. At each step, the value function is approximated using feedforward neural networks. This choice for the numerical approximator was guided primarily by the dimensionality of the problem. In particular, since our problem involves six state variables, we needed a functional approximator that is less subject to the curse of dimensionality than the more commonly used approximation schemes (for instance, polynomials and splines). While increasing the dimensionality of the problem does involve some adjustments, our method can straightforwardly be adjusted to allow for additional dimensions. A discussion of neural networks and the details of our numerical procedure are presented in Appendix A.

In the two following subsections, we present the functional forms and parameters used when the model is solved numerically. It is important to note that, except where explicitly stated, qualitatively similar results could be obtained by using different functional forms and parameters. In the third subsection, we present our numerical results. The parameter values and starting values for the state variables are summarized in Table 2.1. Table 2.2 shows selected results from our numerical analysis.

2.3.1. Chosen functional forms

We assume that the inverse demand curve for the commodity is isoelastic:

$$p_t = P(z_t, S_t) = \frac{z_t}{S_t^{\frac{1}{\epsilon}}} \tag{2.16}$$

with $\epsilon > 1$.

The production function is linear in the capital stock and production shock. That is, $f(\varepsilon_t, K_t) = \varepsilon_t K_t$. Irreversible investment in capital stock is possible, but costly. In order to increase the level of capital stock, the agent must incur adjustment costs that depends both on the amount invested and on the capital stock in place. More specifically, the adjustment cost function is given by:

$$a(I_t, K_t) = \begin{cases} cI_t + b\left(\frac{I_t}{K_t}\right)^2 K_t & \text{if } I_t > 0 \\ 0 & \text{if } I_t = 0 \\ \infty & \text{if } I_t < 0. \end{cases}$$

Due to the quadratic cost component, investment in capital stock benefits the agent not only by increasing future output, but also by reducing future adjustment costs.

For simplicity, we assume that the production and consumption shocks have *iid* binomial distributions, and that the shocks are independent of each other. As will be discussed below, qualitatively similar results could be obtained if we assumed different stochastic processes for the production and consumption shocks (such as mean-reverting processes), as long as the shocks are temporary in nature.

We assume that the demand shock is deterministic and follows a sinusoidal pattern:

$$z_{t+1} = 0.5\left(\sin\left(\xi_{t+1}\right) + 1.5\right) \text{ with } \xi_{t+1} = \xi_t + \vartheta\pi.$$

This form for the demand shock allows us to explore how the storage decision is affected by an upward or downward trend in demand. The number of periods required to complete a cycle depends on the value of the parameter ϑ. We set ϑ equal to 2 to examine the case with no demand shock.

For simplicity, we assume that the agent's utility function is given by

$$u(C_t, \varsigma_t) = \frac{1}{1-\gamma}\varsigma_t C_t^{1-\gamma}$$

with $\gamma > 0$.

2.3.2. Parameters

To solve the model numerically, we must choose parameters values. Unless otherwise specified, the following parameters will be held constant throughout our analysis:

- The subjective discount rate (β) is 0.9.
- There are adjustment costs associated with an increase in the level of the capital stock. More specifically, we have

$$a(I_t, K_t) = \begin{cases} 0.6I_t + 0.2\left(\frac{I_t}{K_t}\right)^2 K_t & \text{if } I_t > 0 \\ 0 & \text{if } I_t = 0 \\ \infty & \text{if } I_t < 0. \end{cases}$$

The specific numbers chosen for the linear and quadratic components of the adjustment cost function are such that investment in capital stock is cheap enough that the agent can respond to shocks through investment, but costly enough that storage impacts the longer term price dynamics.

- The depreciation rate of the capital stock (δ) is 0.2. The depreciation rate was obtained using a balanced panel of firms in the gold and silver mining industry (SIC codes 1040 and 1041), from 1986 to 2004, in Compustat. For each year t, for each firm i, we calculate the average useful life of the capital stock and then calculate the depreciation rate as 2 divided by the median of the useful life for all firms.
- The production and consumption shocks are equally likely to be either 1.5 or 0.5.
- The length of the demand cycle is 8 periods (i.e. $\vartheta = 0.25$). That is, it takes 8 periods for the demand shock to revert to its starting point.
- The coefficient of elasticity (ϵ) is 2.
- The curvature of the utility function (γ) is 0.5. Note that increasing γ makes our results stronger.

2.3.3. Numerical results

In order to study how a change in the importance of the financial role, relative to the industrial role, of the commodity affects the price dynamics, our numerical analysis focuses on the spot price dynamics when the stock of liquidity is initially low. The initial limited

Table 2.1: **Summary of parameter values and starting values for the numerical analysis.**

Parameter or state variable	Value
β	0.9
p	0.6
b	0.2
δ	0.2
ϑ	0.25
ϵ	2.0
γ	0.5
W_0	0.0
Q_0	0.0
K_0	6.0
ε_0	0.5
z_0	$-\pi/2$
ς_0	1.5

availability of liquidity will give the commodity a greater financial role in earlier periods.[7] As time passes, the stock of the liquid asset will increase, which will reduce the importance of the financial role played by the commodity.

We choose starting values which reflect the initial limited availability of liquidity. We assume that the agent initially has no liquid wealth and no inventory of the commodity ($W_0 = Q_0 = 0$), that the production shock is initially low ($\varepsilon_0 = 0.5$) and that the demand shock starts off at its lowest level $\left(\xi_0 = -\frac{\pi}{2}\right)$. To reflect the high demand for liquidity, we assume that the consumption shock is initially high ($\varsigma_0 = 1.5$). Finally, we set the initial level of capital stock to $K_0 = 6$.

[7]We note that, throughout our analysis, when we refer to periods during which the commodity has a more prominent financial role, we are referring to periods during which there is a higher dependence of the storage decision on the demand for liquidity.

Since the computation of the n-period (prepaid) forward price involves taking an expectation over all possible states of the world in n periods (see equation (2.14)), it is difficult to track the effect of the various drivers of the spot price at different points in time by examining the forward curve. For this reason, the first part of our numerical analysis focuses on the statistical properties of our model. In particular, we study the effect of the supply of liquidity on the joint dynamics of the spot price and the convenience yield over time.

The second part of our numerical analysis looks at the conditional variance of the forward price and at the expected level of investment in capital stock.

2.3.3.1. Spot price and convenience yield. Initially, since the agent is not endowed with any liquid wealth, the storage decision for the commodity is affected by the agent's short term consumption needs. Since the consumption shock is high, the agent's demand for the liquid asset is high such that the current consumption value of the commodity exceeds its future asset value. This gives rise to a convenience yield. As time passes and the demand shock for the commodity increases, the agent is able to start accumulating more liquid wealth and a smaller portion of the current sale proceeds are necessary to meet the agent's short term consumption needs. Put differently, as time passes, the storage decision for the commodity depends less on the consumption decision of the agent, and more on the level of the commodity inventory, and on the production and demand shocks. As such, the agent has more incentive to carry some inventory of the commodity into the next period, to take advantage of the higher demand for the commodity in the next period. The decreased need for liquidity thus puts downward pressure on the convenience yield, and the lower quantity supplied puts upward pressure on the spot price (see equation

(2.16)). As more liquid wealth is accumulated and the storage decision becomes less dependent on the consumption decision, the likelihood that a rise in the spot price is accompanied by a fall in the convenience yield falls.

To see why a decrease in the demand for liquidity puts downward pressure on the convenience yield, consider the expression for the one-period convenience yield:

$$q_{t,t+1} = \frac{p_t u'(C_t)}{p_t u'(C_t) - \lambda_t} - 1. \tag{2.17}$$

Now, re-arranging equation (2.7), we have that:

$$\lambda_t = p_t \beta E_t [V_W(\Lambda_{t+1})] - \beta E_t [V_Q(\Lambda_{t+1})] + \mu_t p_t. \tag{2.18}$$

Combining equations (2.17) and (2.18) shows that the one-period convenience yield is monotonically increasing in the Lagrangian multiplier for the liquidity constraint (μ_t). Since μ_t represents the shadow value of consuming an additional unit of liquid wealth at time t, this gives us the result that the one-period convenience yield depends, at least partially, on the value associated with having liquidity at a particular point in time. And since the value placed on liquidity increases with the demand for liquidity when the stock of liquid wealth is low, the one-period convenience yield increases with consumption when liquidity is tight.

We now examine our numerical results more closely, focusing on the likelihood that a rise in the spot price is accompanied by a fall in the convenience yield. As Table 2.2 shows, given our initial starting values, a rise in the spot price is initially accompanied by a fall in the convenience yield if the agent is less liquidity constrained in the second period (i.e. at $t = 1$).[8] This means that there is a 75 percent probability that a rise in the spot price will be accompanied by a fall in the convenience yield as we move from period one to period two.

The intuition for the spot price rising and the convenience yield falling when the agent moves from a state where he is highly liquidity constrained to a state where he is less liquidity constrained is as follows. An agent who is highly liquidity constrained values the liquidity provided by the commodity and, as such, sells more of the commodity than if he were not liquidity constrained. When the agent becomes less liquidity constrained, ceteris paribus, the quantity sold falls, which pushes the price up due to the downward sloping demand curve. In addition, the reduced need for liquidity puts downward pressure on the convenience yield since the option to sell the commodity in the current period which comes from physical ownership of the commodity is not as highly valued.

Two observations about the results presented in Table 2.2 are noteworthy. First, for the three (out of four) cases when the agent is less liquidity constrained in the second period than in the first, three quantities are of particular interest: the level of liquid wealth carried into the third period ($W_2 = 0$), the level of commodity inventory carried into the third period ($Q_2 > 0$), and the level of investment in capital stock during the second period ($I_1 = 0$). Initially, the agent does not have any liquid wealth and must sell his inventory of the commodity in order to consume. Because his need for liquidity is initially high and his stock of the commodity is low, the agent initially stocks out. If the agent is less liquidity constrained in the following period, then the agent can meet his

[8]The agent becomes more liquidity constrained as he moves from the first to the second period only if the production shock in the second period is again low and the consumption shock is again high, due to the depreciation of the capital stock.

consumption needs by selling only a part of the commodity on hand and carrying over the rest into the next period. Also, regardless of the state of the world in the following period, the agent does not invest in capital stock in the second period. This result is most intuitive when thought of in a portfolio allocation framework. Recall that the agent can move wealth through time by storing the commodity, investing in capital stock, and accumulating liquid wealth. Our results can thus be interpreted as the agent choosing to move wealth through time using only lower risk mechanisms when liquidity is tight, due to his desire to smooth consumption.[9] We note that the reason that the agent sells no more of the commodity than needed to meet his consumption needs in the second period is the upward trend in the demand shock. If the demand shock were constant over time, then the agent would move more of his wealth through time using the liquid wealth account.

Second, when the agent is less liquidity constrained in the second period, the convenience yield falls to zero. Since the convenience yield is bounded below by zero, this means that subsequent moves in the spot price cannot be accompanied by falls in the convenience yield. This is due to the binomial nature of the production and liquidity shocks. If we had chosen a process with a greater number of possible shock magnitudes in each period (as opposed to a process where only high or low shocks are possible), then the negative relation between the spot price and the convenience yield could persist for more than one period. The idea is that if there are intermediate shock levels, then the agent may move more slowly out of his liquidity constraint, which would be reflected by a convenience yield more slowly moving towards zero. We note that the negative relation between the spot price and the convenience yield could also persist for more than one

[9] Recall that the demand shock is deterministic.

period if the demand shock were stochastic, due to the greater use of the liquid wealth account (and the reduced use of the storage of the commodity) to carry wealth over to the next period if storage is riskier.

If the cost of adjusting the level of the capital stock is increased, then storage plays a more prominent role due to the depreciation of the capital stock.[10] Table 2.3 compares the probability that a rise in the spot price will not be accompanied by a rise in the convenience yield when investment is and is not possible.[11] As can be seen from this table, investment in capital stock makes a rise in the spot price less likely to not be accompanied by a rise in the convenience yield for an extended period of time. In the case of infinite adjustment costs, a rise in the spot price must be accompanied by a rise in the convenience yield after three periods if the production shock is low in all periods. This is because, with the current model setup, if the production shock is consistently low, then the quantity available for sale in a given period will eventually be so low (due to the depreciation of the capital stock) that the convenience yield will eventually start rising.

Since it may appear as though our results are driven by the special form chosen for the demand shock process, it is important to note that a rise in the spot price can be accompanied by a fall in the convenience yield even in the extreme case that the demand shock is constant (i.e. $\vartheta = 2$). This relation, however, is less likely to occur. The reduced likelihood is due to a lower probability that the spot price rises since, with a constant demand shock, there is no increase in the exogenous demand shock to counterbalance the

[10] Recall that we assume that the commodity inventory does not depreciate.

[11] Note that the first probability in each row represents the probability that a rise in the spot price is accompanied by a fall in the convenience yield.

reduction in the price that comes from an increase in the quantity sold when the agent becomes less liquidity constrained due to an increase in the production shock.

If we increase the degree of risk aversion, then a rise in the spot price is more likely to not be accompanied by a rise in the convenience yield, and this relation is more likely to persist. The higher likelihood is due to the fact that an increase in the degree of risk aversion represents an increase in the agent's desire to smooth consumption over time. The greater desire to smooth consumption is reflected by a greater sensitivity of the convenience yield to the demand shock when the liquidity constraint is binding. That is, the convenience yield is more likely to fall as the proceeds from the sale of the commodity increase. The higher persistence is explained by the fact that, with a higher degree of risk aversion, it will take the agent more time to accumulate enough liquid wealth to significantly reduce his sensitivity to the timing of the proceeds from the sale of the commodity.

Our analysis of the joint dynamics of the spot price and convenience yield is not very sensitive to the level of the starting value for the capital stock. Our results are also qualitatively similar if we vary the beginning inventory of the commodity, as long as the inventory is low relative to the capital stock in place. Finally, we note that our results are not due to the specification for the production and consumption shocks. Qualitatively similar results could be obtained using different processes for the production and consumption shocks, as long as these shocks are temporary in nature.

Our numerical results show that a rise in the spot price can be accompanied by a fall in the convenience yield when the agent becomes less liquidity constrained over time. As might be expected, a fall in the spot price can be accompanied by a rise in the convenience yield if the agent becomes more liquidity constrained over time. The analysis of this case (not presented here) shows that such a relation cannot be sustained for more than two periods. The reason for this is that, when the stock of liquid wealth is low, the level of investment in capital stock (if any) is low, such that the incentive to store the commodity is higher, due to the depreciation of the capital stock. The heightened incentive to store the commodity puts downward pressure on the convenience yield and reduces the likelihood that the convenience yield rises over time.[12]

In summary, our numerical results show that the commodity spot price and the convenience yield can move in opposite directions when the stock of liquidity is low. This result is driven by the fact that, if the stock of liquidity is low, then the agent's commodity storage decision depends, at least partially, on his demand for liquidity. If more liquid wealth is available, then the agent is less sensitive to the timing of the cash flows and, in this case, the spot price and the convenience yield are less likely to move in opposite directions.

Table 2.2: **Selected results.** Selected results from the numerical simulation at $t = 0$, $t = 1$ and $t = 2$ when $W_0 = 0$, $Q_0 = 0$, $K_0 = 6$, $\varepsilon_0 = 0.5$, $\varsigma_0 = 1.5$, and $\xi_0 = -\frac{\pi}{2}$.

$t=0:$	$p_0 = 0.14$	$q_{0,1} = 0.15$										
$t=1,2:$												
ε_1	ς_1	p_1	$q_{1,2}$	Q_1	K_1	W_1	S_1	I_1	C_1	Q_2	K_2	W_2
High	High	0.15	0.00	0.00	4.8	0.00	6.55	0.00	1.01	0.65	3.84	0.00
High	Low	0.30	0.00	0.00	4.8	0.00	1.80	0.00	0.53	5.40	3.84	0.00
Low	High	0.26	0.15	0.00	4.8	0.00	2.40	0.00	0.61	0.00	3.84	0.00
Low	Low	0.32	0.00	0.00	4.8	0.00	1.51	0.00	0.49	0.89	3.84	0.00

[12]Since the effect of the depreciation of the capital stock on the incentive to store increases the likelihood that the convenience yield falls (or, at least, does not increase) over time, the depreciation of the capital stock increased the likelihood that a rise in the spot price would be accompanied by a fall in the convenience yield for more than one period in the case studied above.

Table 2.3: **Effect of adjustment cost.** This table shows the probability that a rise in the spot price is not accompanied by a rise in the convenience yield from time t to time $t+1$, conditional on a rise in the spot price being accompanied by a fall in the convenience yield in all previous periods, when $W_0 = 0$, $Q_0 = 0$, $K_0 = 6$, $\varepsilon_0 = 0.5$, $\varsigma_0 = 1.5$, and $\xi_0 = -\frac{\pi}{2}$. The first column indicates the beginning of the period under consideration. The column labeled "Investment" shows the probability when $p = 0.6$ and $b = 0.2$ in the adjustment cost function. The column labeled "No investment" shows the probability when adjustment costs are infinite.

t	Investment	No investment
0	0.75	0.75
1	0.25	0.38
2	0.08	0.19
3	0.00	0.06

Table 2.4: **Effect of quadratic adjustment cost on investment.** This table shows the expected level of investment at time t when $W_0 = 30$, $Q_0 = 0$, $\varepsilon_0 = 0.5$, $\varsigma_0 = 1.5$, and $\xi_0 = -\frac{\pi}{2}$. The second and third columns show the results for $K_0 = 1$ and $K_0 = 6$, respectively.

	K_0	
t	1	6
0	1.20	2.34
1	1.68	1.58
2	1.84	1.23
3	1.53	1.13
4	0.81	0.88

2.3.3.2. Conditional variance of forward prices and investment. Figure 2.4 summarizes the numerical results for the conditional variance of the forward price and the investment in capital stock (the results for the case with $W_0 = 30$ are shown in Figure

2.5 for comparison). The top plot shows the expected investment in capital stock at various horizons. The bottom plot shows the conditional variance of the forward price for contracts of different maturities.

Consistent with Routledge, Seppi, and Spatt (2000) and Kogan, Livdan, and Yaron (2005), we find that conditioning on the initial state variables affects the variability of forward prices at various horizons. In particular, we find that the variance of the forward price is affected by the probability (and severity) of a future stockout. We note that the high variance of the three period forward contract is due to our specification for the demand shock.

Figure 2.4 shows that, when the initial stock of liquidity is low, the investment in capital stock depends on the availability of liquid wealth and is highly sensitive to the path of the demand shock.[13] Figure 2.5 shows that, if the initial endowment in liquid wealth is high, then the investment in capital stock is smoother over time due to the quadratic component of the adjustment cost. Since the quadratic component makes the cost of investing lower when the capital stock in place is higher, the effect of the quadratic cost can also be seen by comparing the investment in capital stock for different levels of beginning capital stock. Table 2.4 compares the investment in capital stock when $K_0 = 1$ and $K_0 = 6$, assuming $W_0 = 30$. As can be seen from this table, when there is little capital stock in place, it is optimal to initially invest increasingly over time due to the lower adjustment cost associated with a higher level of capital stock in place.

[13] As such, the fall in investment in period four is due to the fall in the demand shock in period five (recall that there is a lag between the time of investment and the productive use of the new capital stock).

Figure 2.4: Expected level of investment and conditional variance of the forward price at various horizons when $W_0 = 0$. This figure shows the expected level of investment and the conditional variance of the forward price at various horizons when $W_0 = 0$, $Q_0 = 0$, $K_0 = 6$, $\varepsilon_0 = 0.5$, $\varsigma_0 = 1.5$, and $\xi_0 = -\frac{\pi}{2}$. The conditional variance of the forward price is computed as $var\left(F_{t+1,t+1+n} \mid \Lambda_t\right) = E_t\left[F^2_{t+1,t+1+n}\right] - E_t\left[F_{t+1,t+1+n}\right]^2$, with $n = 0, 1, 2, 3$ ($n = 0$ represents the conditional variance of the spot price

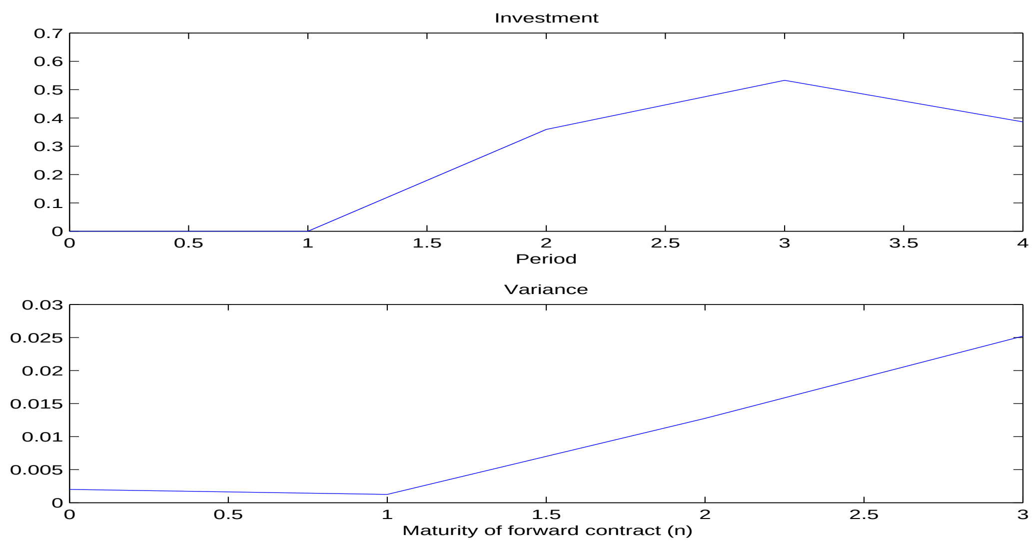

2.4. Empirical analysis

The numerical analysis presented in the previous section suggests that the supply of liquidity can affect the price dynamics of a commodity that serves both as an industrial commodity and a financial asset. More specifically, our numerical results suggest that the convenience yield of commodity that can serve as a financial asset depends, at least partially, on the value placed on liquidity. This is the prediction of the model tested in this section.

Our empirical analysis will be done using data for gold, silver, crude oil and copper. Since gold and silver serve both as industrial commodities and financial assets, our model predicts that the convenience yield of gold and silver is positively related to the value

Figure 2.5: Expected level of investment and conditional variance of the forward price at various horizons when $W_0 = 30$. This figure shows the expected level of investment and the conditional variance of the forward price at various horizons when $W_0 = 30$, $Q_0 = 0$, $K_0 = 6$, $\varepsilon_0 = 0.5$, $\varsigma_0 = 1.5$, and $\xi_0 = -\frac{\pi}{2}$. The conditional variance of the forward price is computed as $var\left(F_{t+1,t+1+n} \mid \Lambda_t\right) = E_t\left[F^2_{t+1,t+1+n}\right] - E_t\left[F_{t+1,t+1+n}\right]^2$, with $n = 0, 1, 2, 3$ ($n = 0$ represents the conditional variance of the spot price).

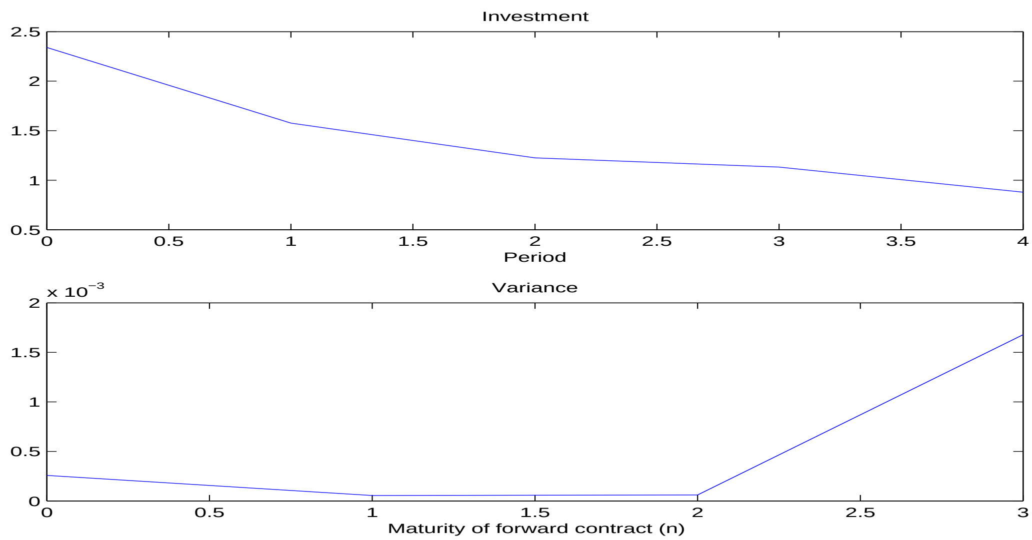

placed on liquidity. The model predicts no positive association between the value placed on liquidity and the convenience yield of oil or copper.[14]

Before presenting our empirical results, we describe our data. We then present a preliminary empirical examination of the relation between the value of liquidity and the convenience yield of the four commodities. In the last subsection, we present our regression analysis.

[14]The reason why gold and silver, as opposed to crude oil and copper, are used as financial assets is not clear. It could be due to their physical properties. For instance, both metals are dense, which means that they are cheap to store. Also, both metals are highly durable, which means that there is no (or little) spoilage associated with the storage these commodity. Perhaps, it is due to their historical association with the monetary system. While there are other commodities which could potentially serve as financial assets (for instance, copper is almost as dense as silver and is durable), the bottom line seems to be that investors do not use these other industrial commodities as financial assets and, as such, these other commodities are not priced as financial assets.

2.4.1. Description of the data

The convenience yield is obtained using futures contracts traded on NYMEX, with the spot price defined as the futures price of the nearest-to-maturity contract. The n period convenience yield at time t is computed as $q_{t,t+n} = \left(\frac{F_{t,t+\varepsilon}}{F_{t,t+\varepsilon+n}}(1 + r_n n) - 1\right)/n$ where $F_{t,t+\varepsilon}$ is the nearest-to-maturity futures contract, $F_{t,t+\varepsilon+n}$ is the n-period futures contract, and r_n is the n-period constant-maturity Treasury rate. Data for futures prices and Treasury rates are obtained from Datastream.

Recent work in the credit risk literature suggests that the variation in the AAA-Treasury yield spread (the "corporate spread") cannot be explained by the variation in the default risk (for instance, Chen, Collin-Dufresne, and Goldstein (2004)).[15] Additionally, the results of Krishnamurthy and Vissing-Jorgensen (2006) suggest that that the variation of the corporate spread is mostly due to variations in the "convenience yield"of Treasuries. In particular, Krishnamurthy and Vissing-Jorgensen find that there is a negative relation between the corporate spread and the stock of debt, measured as the ratio of the privately held Treasury debt relative to GDP. Motivated by these results, we use the yield spread between Moody's AAA long maturity bonds and 10-year Treasuries as our measure of the value placed on liquidity.[16] AAA and 10-year Treasury yields from 1983 to 2005 are obtained from the website of the Federal Reserves Bank of St-Louis. In the analysis below, we refer to this measure as the "corporate spread".

[15]More precisely, the authors find that the variation in the historical default risk is too low to explain the variation in the corporate spread

[16]Krishnamurthy and Vissing-Jorgensen (2006) use the difference between Moody's AAA long maturity bonds and long term maturity (>10 years) Treasuries. We use 10-year Treasuries instead to obtain a larger sample. We note that, adjusting the sample period, qualitatively similar results can be obtained using longer maturity Treasuries.

We use the Debt to GDP ratio to create a dummy for the supply of liquidity. This ratio is constructed as the ratio of privately held Treasury debt relative to GDP, where private debt includes debt held by Federal Reserve, but excludes debt held by other parts of the government such as the Social Security Trust Fund. Data up to 2003 was obtained from Henning Bohn's web page. Data for 2004 and 2005 was obtained from the Office of Management and Budget.

To control for the changes in the default risk, we use the difference between Moddy's BAA and AAA long maturity bond yields since we expect this spread to widen when default risk increases. BAA yields from 1986 to 2005 are obtained from the website of the Federal Reserves Bank of St-Louis. In the analysis below, we refer to this measure as the "credit spread".

Finally, we control for the slope of the yield curve, defined as the difference between the yield on 10-year Treasury Bonds and 3-month Treasury Bills. The slope of the yield curve is used to control for the risk premium due to economic conditions. The data for the 3-month Treasury bill yields are from the website of the Federal Reserves Bank of St-Louis. In the analysis below, we refer to this measure as the "slope of the term structure".

2.4.2. Preliminary results

The main implication of our model is that the convenience yield of gold and silver is partially determined by the value placed on liquidity. As such, we expect to see a positive association between the corporate spread (which proxies for the value placed on liquidity) and the one-period convenience yield of gold and silver. Furthermore, we expect the convenience yield of gold and silver to be more sensitive to the value placed on liquidity when the supply of liquidity is low. We thus expect the positive association between the corporate spread and the convenience yield to be stronger in periods when the Debt to GDP ratio is low.

Figures 2.6 to 2.9 show scatter plots of the corporate spread against the one-year convenience yield of gold, silver, crude oil and silver. The top plots show the scatter plots for years during which the Debt to GDP ratio was below 0.39 (i.e. years 1983 to 1985 and 2000 to 2005). The bottom plots shows the scatter plots for years during which the Debt to GDP ratio was at least 0.39 (i.e. years i.e. 1986 to 1999). The scatter plots were obtained using daily data from 1983 to 2005 for gold, silver and oil, and from 1989 to 2005 for copper.

Figure 2.6 shows that, as expected, there appears to be a positive association between the convenience yield and the corporate spread, particularly when the stock of liquidity is low. For silver, there appears to be only a very slight positive association between the corporate spread and the convenience yield.[17]

Figures 2.8 and 2.9 show that, for oil and copper, there does not appear to be a positive relation between the convenience yield and the corporate spread. In fact, there appears to be a negative relation. This negative relation can be understood in the context of the model as follows. Recall that when liquidity is highly valued, the agent chooses to invest only in the lower risk technologies. As such, he transfers wealth through time by storing the commodity and using the wealth account, but does not invest in capital stock. The greater use of storage and the lack of investment in capital stock puts downward pressure

[17]That the positive relation between the corporate spread and the convenience yield is weaker for silver is not surprising since silver is more sought after than gold in industrial uses.

Figure 2.6: Corporate spread against one-year convenience yield of gold.

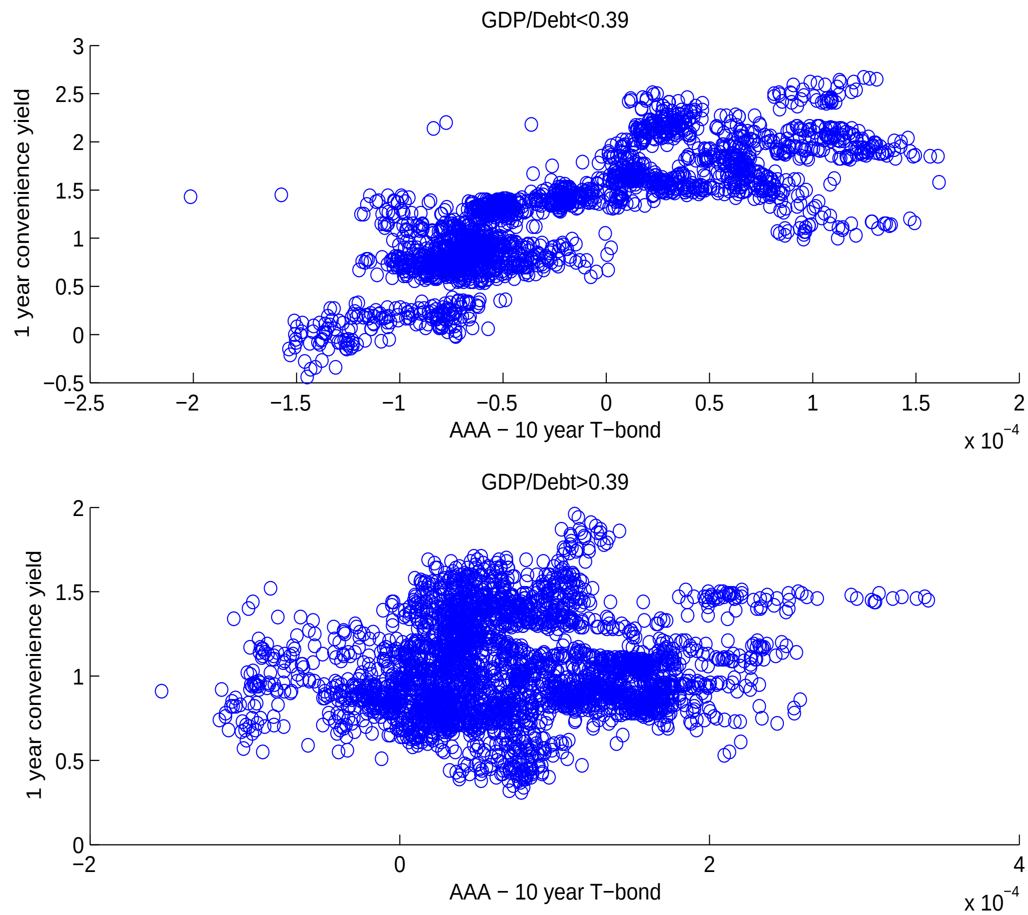

on the convenience yield. The difference with the case of gold and silver comes from the fact that gold and silver can be immediately converted into liquid wealth. The greater need for liquidity thus leads the agent to sell more gold or silver, which puts upward pressure on the convenience yield.

Figure 2.7: Corporate spread against one-year convenience yield of silver.

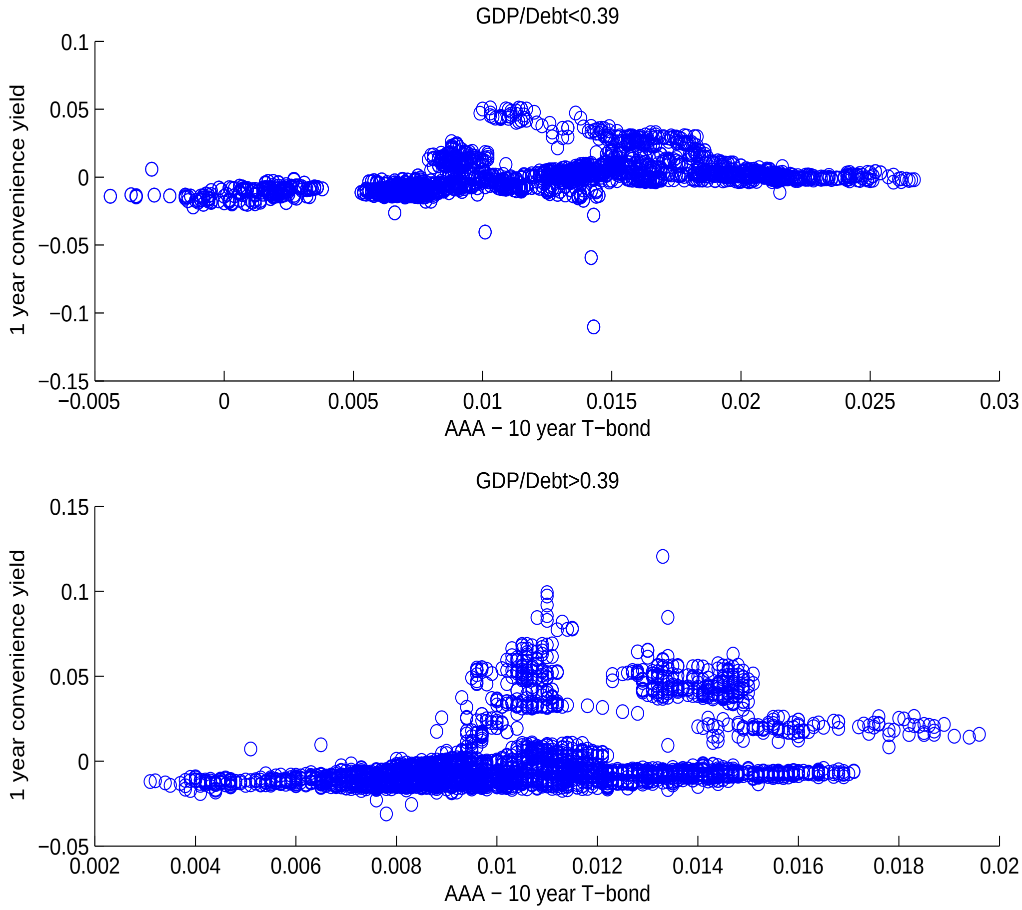

2.4.3. Regression analysis

To examine more rigourously whether or not the value placed on liquidity affects the convenience yield of gold, silver, crude oil and copper, we now turn to regression analysis. The regression analysis was done using monthly data from 1986 to 2005 for gold, silver and oil, and from 1989 to 2005 for copper. Similar results were obtained using daily data.

Figure 2.8: Corporate spread against one-year convenience yield of crude oil.

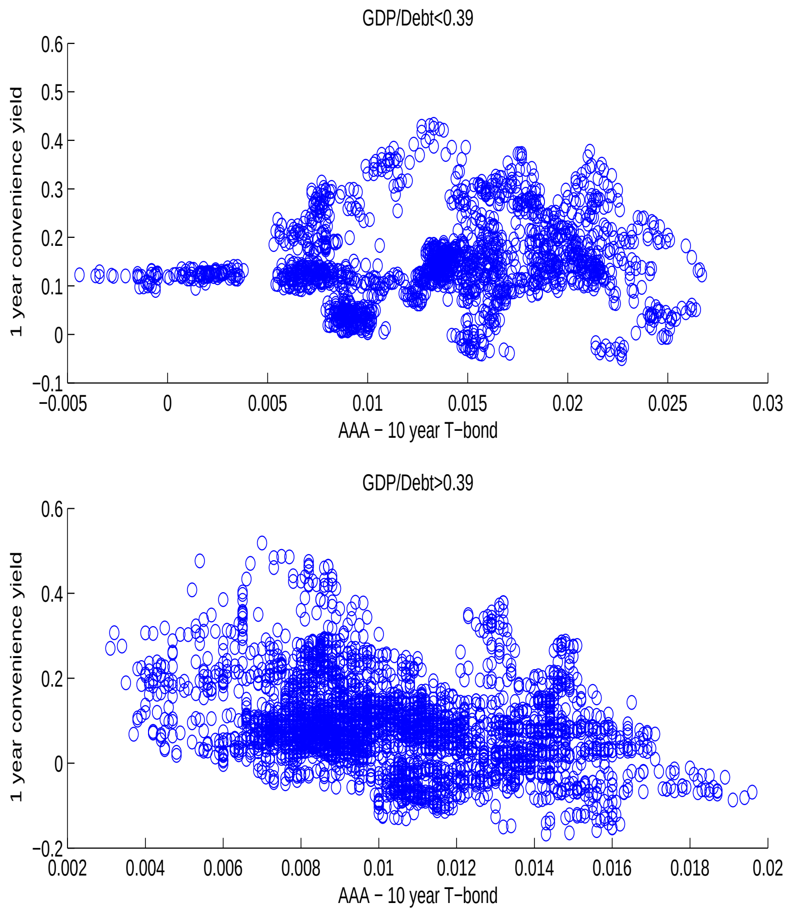

We run the following regression:

$$q_{t,t+12} = \alpha + \beta x_{t-1} + \varepsilon_t$$

where $q_{t,t+12}$ is the one-year convenience yield and x_{t-1} represents the independent variables, lagged one month. The independent variables are:

Figure 2.9: Corporate spread against one-year convenience yield of copper.

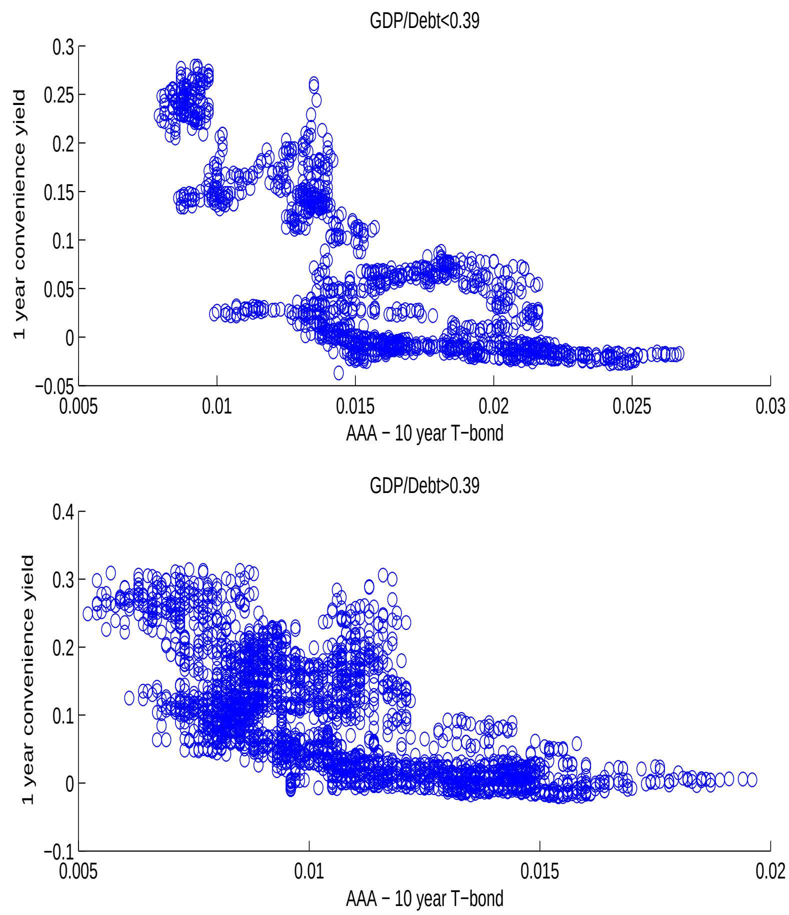

- low_debt_to_GDP$_t$: a dummy variable equal to one if the Debt to GDP ratio was below 0.39 during the year of period t. This dummy variable is used to control for high and low stocks of liquidity.
- $q_{t-1,t+11}$: the lagged one-year convenience yield.
- $AAA_{t-1} - r_{10,t-1}$: the lagged corporate spread.
- $r_{10,t-1} - r_{0.25,t-1}$: the lagged slope of the term structure.
- $BAA_{t-1} - AAA_{t-1}$: the lagged credit spread.

- $S_{0,t-1}$: the lagged spot price.
- res_{t-1}: the lagged residuals of the regression of the spot price on the corporate spread and a constant. Using these residuals allows us to control for the portion of the spot price which cannot be explained by changes in the corporate spread. Using these residuals instead of the lagged spot price affects our regression results due to the high correlation between the corporate spread and the spot price.

Tables 2.5 to 2.8 show our regression results for gold, silver, crude oil and copper, respectively.

The results of the regression analysis for gold (Table 2.5) show that there appears to be a positive relation between the convenience yield and the corporate spread. The coefficient on the corporate spread is positive and significant, at the 10 percent level once we control for the slope of the term structure, and at the 5 percent level once when we also control for the credit spread. The reason that the coefficient on the corporate spread increases (both in magnitude and in significance) when we control for the slope of the term structure and the credit spread could be that, without these controls, the corporate spread acts as a proxy for economic conditions, rather than a proxy for the value placed on liquidity.[18] Put differently, the increase in the significance of the coefficient on the corporate spread once we control for economic conditions could be because, for gold, it is important to distinguish between the effect of economic conditions on the production process (for instance, investment in capital stock may be lower in bad times), and the effect of demand for liquidity on the storage decision. For a commodity which cannot

[18]Fama and French (1988) show that the difference between the average yield of bonds rated BAA and AAA by Moody's tracks long term business cycles, and that the difference between the average yield of Treasury bonds with more than 10 years to maturity and the yield of 3-month Treasury bills tracks the short-term business cycles.

provide liquidity, making this distinction may not be as important. Our regression results are consistent with this. To see this, consider the coefficients on the corporate spread, the slope of the term structure, and the credit spread. In the regressions for silver, oil and copper, at most one of these coefficients is significant.

For silver (see Tables 2.6), the coefficient on the corporate spread is positive, which is consistent with the prediction of our model, but this coefficient is not statistically significant. For crude oil and copper (Tables 2.7 and 2.8), the coefficient on the corporate spread is negative. This coefficient is only significant in the case of copper. The negative coefficient can be interpreted as oil and copper producing firms making greater use of storage and investing less in capital stock when liquidity is tight. For gold and silver, the greater need for liquidity could be leading producers to sell more of the commodity, which puts upward pressure on the convenience yield.

Before concluding, we comment on the negative sign of the dummy variable low_debt_to_GDP_t for gold (see Table 2.5). In particular, we highlight why the negative sign on this dummy does not contradict the predictions of the model. For simplicity, let us assume for now that the value placed on liquidity is zero, such that the liquidity component of the convenience yield is negligible. If the market stock of liquidity is low, then it is reasonable to expect that the industrial demand for the commodity will be low. If this is the case, then storers of gold may decide to store more of the commodity, to wait for more favorable selling conditions. The increased use of storage would put downward pressure on the convenience yield. Loosely speaking, if we expect the industrial demand for the commodity to increase with the stock of liquidity, then we expect the convenience yield to be higher when the stock of liquidity is high.

Overall, our regression results show some support for the prediction of our model that the convenience yield of gold and silver is positively related to the value placed on liquidity. This, of course, assumes that the corporate spread is a good proxy for the value placed on liquidity.

Table 2.5: **Regression results for gold.** The dependent variable is the one-year convenience yield of gold. Reported are the coefficients and the absolute value of the robust t-statistics (in parenthesis) for the independent variables. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
low_debt_to_GDP$_t$	-0.00096	-0.00133	-0.00175	-0.00158	-0.00179	-0.00190	-0.00190
	(2.67)***	(2.90)***	(4.38)***	(3.31)***	(3.47)***	(3.79)***	(3.79)***
$q_{t-1,t+11}$	0.88438	0.87610	0.81145	0.83366	0.81401	0.76590	0.76590
	(19.05)***	(18.59)***	(14.97)***	(15.84)***	(14.10)***	(12.41)***	(12.41)***
$AAA_{t-1} - r_{10,t-1}$		0.00059		0.00089	0.00120	0.00017	0.00120
		(1.13)		(1.69)*	(2.00)**	(0.28)	(2.05)**
$S_{0,t-1}$			-0.00002			-0.00001	
			(3.15)***			(2.71)***	
$r_{10,t-1} - r_{0.25,t-1}$				-0.00050	-0.00048	-0.00041	-0.00041
				(2.83)***	(2.75)***	(2.47)**	(2.47)**
$BAA_{t-1} - AAA_{t-1}$					-0.00158	-0.00120	-0.00120
					(1.65)*	(1.35)	(1.35)
res_{t-1}							-0.00001
							(2.71)***
α	0.00093	0.00038	0.00716	0.00116	0.00234	0.00877	0.00220
	(3.12)***	(0.61)	(3.57)***	(1.57)	(2.33)**	(3.23)***	(2.31)**
Observations	238	238	238	238	237	237	237
Adjusted R^2	0.82	0.82	0.83	0.83	0.83	0.83	0.83

Table 2.6: **Regression results for silver.** The dependent variable is the one-year convenience yield of silver. Reported are the coefficients and the absolute value of the robust t-statistics (in parenthesis) for the independent variables. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
low_debt_to_GDP$_t$	-0.00029	-0.00095	-0.00022	-0.00072	-0.00072	-0.00070	-0.00070
	(0.27)	(0.67)	(0.20)	(0.50)	(0.51)	(0.48)	(0.48)
$q_{t-1,t+11}$	0.93919	0.93535	0.94116	0.91491	0.90204	0.90241	0.90241
	(19.06)***	(19.00)***	(19.18)***	(17.55)***	(16.58)***	(16.65)***	(16.65)***
$AAA_{t-1} - r_{10,t-1}$		0.00133		0.00144	0.00186	0.00182	0.00184
		(0.95)		(1.03)	(1.36)	(1.30)	(1.34)
$S_{0,t-1}$			-0.00000			-0.00000	
			(0.93)			(0.09)	
$r_{10,t-1} - r_{0.25,t-1}$				-0.00062	-0.00058	-0.00058	-0.00058
				(1.65)*	(1.55)	(1.56)	(1.56)
$BAA_{t-1} - AAA_{t-1}$					-0.00268	-0.00266	-0.00266
					(1.83)*	(1.79)*	(1.79)*
res_{t-1}							-0.00000
							(0.09)
α	0.00030	-0.00110	0.00191	-0.00015	0.00168	0.00185	0.00168
	(0.63)	(0.74)	(1.09)	(0.09)	(0.82)	(0.73)	(0.82)
Observations	220	220	220	220	219	219	219
Adjusted R^2	0.87	0.87	0.87	0.87	0.87	0.87	0.87

Table 2.7: **Regression results for crude oil.** The dependent variable is the one-year convenience yield of crude oil. Reported are the coefficients and the absolute value of the robust t-statistics (in parenthesis) for the independent variables. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
low_debt_to_GDP$_t$	0.00841	0.01303	0.02394	0.01473	0.01568	0.05298	0.05298
	(0.95)	(1.31)	(1.79)*	(1.45)	(1.55)	(2.78)***	(2.78)***
$q_{t-1,t+11}$	0.79723	0.79359	0.80925	0.78475	0.77470	0.78565	0.78565
	(15.41)***	(15.56)***	(15.88)***	(14.98)***	(14.58)***	(14.99)***	(14.99)***
$AAA_{t-1} - r_{10,t-1}$		-0.00787		-0.00779	-0.00841	-0.02708	-0.03294
		(0.65)		(0.64)	(0.68)	(1.91)*	(2.14)**
$S_{0,t-1}$			-0.00106			-0.00176	
			(2.38)**			(3.06)***	
$r_{10,t-1} - r_{0.25,t-1}$				-0.00385	-0.00469	-0.00469	-0.00469
				(1.25)	(1.47)	(1.52)	(1.52)
$BAA_{t-1} - AAA_{t-1}$					0.01170	0.00635	0.00635
					(0.50)	(0.27)	(0.27)
res_{t-1}							-0.00176
							(3.06)***
α	0.01746	0.02589	0.03655	0.03268	0.02500	0.08138	0.04658
	(3.06)***	(1.94)*	(3.51)***	(2.20)**	(1.19)	(2.57)**	(1.98)**
Observations	231	231	231	231	230	230	230
Adjusted R^2	0.66	0.66	0.66	0.66	0.65	0.66	0.66

Table 2.8: **Regression results for copper.** The dependent variable is the one-year convenience yield of copper. Reported are the coefficients and the absolute value of the robust t-statistics (in parenthesis) for the independent variables. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
low_debt_to_GDP$_t$	0.00237	0.01582	-0.00029	0.01646	0.01656	0.01214	0.01214
	(0.60)	(3.17)***	(0.08)	(3.16)***	(3.09)***	(2.40)**	(2.40)**
$q_{t-1,t+11}$	0.95351	0.87117	0.84555	0.86504	0.86531	0.75950	0.75950
	(28.89)***	(21.67)***	(12.06)***	(19.88)***	(20.01)***	(9.55)***	(9.55)***
$AAA_{t-1} - r_{10,t-1}$		-0.03029		-0.03138	-0.03115	-0.02885	-0.04312
		(4.59)***		(4.34)***	(4.36)***	(4.08)***	(4.25)***
$S_{0,t-1}$			0.00041			0.00041	
			(2.04)**			(1.83)*	
$r_{10,t-1} - r_{0.25,t-1}$				-0.00102	-0.00094	-0.00318	-0.00318
				(0.62)	(0.50)	(1.41)	(1.41)
$BAA_{t-1} - AAA_{t-1}$					-0.00164	0.00655	0.00655
					(0.13)	(0.52)	(0.52)
res_{t-1}							0.00041
							(1.83)*
α	0.00352	0.04330	-0.02763	0.04673	0.04759	0.01104	0.07021
	(1.19)	(4.50)***	(1.82)*	(3.85)***	(3.28)***	(0.49)	(3.47)***
Observations	182	182	182	182	182	182	182
Adjusted R^2	0.88	0.89	0.88	0.89	0.89	0.89	0.89

2.5. Conclusion

In this chapter, we model the price dynamics of a commodity which can serve both as an industrial commodity and a financial asset. The model predicts that the convenience yield of such a commodity is positively related to the value of liquidity. The model does not predict a positive relation between the convenience yield of an industrial commodity (such as copper and oil) and the value placed on liquidity. Our regression results, which use the corporate spread as a proxy for the value placed on liquidity, appear to support this prediction.

Our work represents an important first step to understanding the benefits of including different types of commodities in an investor's portfolio. Also, since our results imply that the nature of a commodity affects the conditional variance of the forward price, our results may have important implications for option pricing.

CHAPTER 3

Does Analyst Disclosure Matter? A Real-Time Analysis (joint with Joseph Engelberg and Jared Williams)

3.1. Introduction

Within the last few years, financial regulators have been particularly concerned about the integrity of analyst recommendations.[1] Several studies have documented a disproportionate number of buy (relative to sell) recommendations (e.g. Elton, Gruber, and Grossman (1986), Stickel (1995) and Malmendier and Shanthikumar (2005)) and others have shown affiliated analysts make optimistic forecasts for current or potential clients (Michaely and Womack (1999), Dugar and Nathan (1995) and Lin and McNichols (1998)). In response, regulators proposed increased disclosures:

- On May 10, 2002, the SEC issued an order that approved rule changes by NYSE and NASD which required analysts who make public appearances to disclose whether they own shares of the recommended securities, whether their firm owns

[1]The following quote by SEC Chairman Arthur Levitt at a town hall meeting in Philadelphia on January 16, 2001, illustrates the concern of regulators: "In many respects, a culture of gamesmanship has taken root in the financial community making it difficult to tell salesmanship from honest advice. And that means, more than ever, investors must stay on their toes. How many of you have seen analysts from Wall Street firms on television talking about one company or another? Many of you probably have not thought twice about that person's recommendation to buy a particular stock. But you should. A lot of analysts work for firms that have business relationships with the same companies these analysts cover. Some analysts' paychecks are tied to the performance of their employers, who make a lot of money underwriting or owning those stocks.... It's your right to know when conflicts exist. And while I can't say it's your obligation, I would ask every investor to help bring pressure to bear on these markets to ensure that greater disclosure is soon a reality."

more than 1% of the outstanding equity of the recommended firm and whether their firm has an investment banking relationship with the recommended firm.[2]

- On July 30th, 2002, President Bush signed into law the Sarbanes Oxley Act of 2002 (SOX). Title V of SOX entitled "Analyst Conflicts of Interest" amends the Securities Exchange Act of 1934 and requires analysts to disclose a financial interest in or association with a firm they review. SOX makes it clear that these disclosures apply to written reports and public appearances and that the disclosure must include any equity and debt holdings the analyst holds in the subject company, any compensation paid to the analyst by the subject company, and whether the subject company is a client of the analyst's firm. A year later, the SEC released a statement confirming that the recently amended rules of the NYSE and the NASD satisfied the requirements of the SOX.[3]

- On April 28, 2003, ten of the largest firms on Wall Street agreed to pay $1.4 billion in penalties in response to government charges that they issued optimistic analyst reports to secure and maintain investment banking clients. In this "global settlement" the 10 firms also agreed to many structural changes including disclosure requirements on written analyst reports.

---

[2] See http://www.sec.gov/rules/sro/34-45908.htm for the official order.

[3] The official release by the SEC on July 29th, 2003 reads: "The SOA requires disclosure of the extent to which a research analyst has debt or equity investments in the issuer that is the subject of the research report or public appearance. Current NASD Rule 2711(h)(1)(A) requires disclosure of whether the 'research analyst or a member of the research analyst's household has a financial interest in the securities of the subject company, and the nature of the financial interest (including, without limitation, whether it consists of any option, right, warrant, future, long or short position).' The Commission believes that NASD Rule 2711(h)(2), and NYSE Rule 472(k)(1) and (2), as amended, satisfy the requirements of Exchange Act 15D(b)(1)." (http://www.sec.gov/rules/sro/34-48252.htm)

While the mandated disclosures do provide information to investors, it is not clear how this information is interpreted. For example, consider the case where an analyst[4] recommends investors buy shares of a certain stock and it is disclosed that the analyst owns shares of this stock. It would be reasonable for investors to think that the recommendation is more credible because the analyst himself is willing to own the stock. It would also be reasonable for investors to discount the recommendation because they are concerned that the analyst is attempting to push up the stock's price and therefore increase his wealth via the recommendation. Finally, it could also be that the existence of laws and penalties for unsubstantiated claims makes disclosure irrelevant since investors now have the sense that claims made by analysts can be corroborated.[5]

The purpose of this paper is to investigate investors' response to ownership disclosure. Our investigation should be useful for several reasons. First, our results should be useful to policymakers who must balance the cost imposed on those who disclose with the benefits individual investors gain from the disclosure.[6] Understanding how investors respond to

[4]Throughout our paper, we use the definition of an "analyst" used in SOX: an associated person of a registered broker or dealer who prepares a written or electronic report with information reasonably sufficient to make an investment decision *irregardless of whether the person is called an "analyst."*

[5]Note that, because the first two interpretations affect price and buy-sell imbalance in opposite directions, it is possible that the disclosure event will appear to have no impact on returns and buy-sell imbalance, even if investors are waiting to see the information contained in the ownership disclosure to make their investment decision. If this is the case, however, we would expect abnormal trading volume following the disclosure event.

[6]In our case of analyst disclosures in public appearances, some media spoke of the additional cost these disclosures would entail. From the Financial Times (July 5, 2002):

Jeff Randall, the BBC business editor, said: "It is going to make life very difficult. In some cases, the disclaimer will be longer than the answer given by the analyst." Mr Randall said, however, that he understood why the regulators were taking this step.

To demonstrate compliance, investment banks will have to keep copies or transcripts of all broadcasts.

Julian Heynes, producer at CNBC, the business television channel, said that one of the biggest problems would be coordination between CNBC's bureaus around the world.

"We are concerned about producing the guidelines in a coherent and logical way," Mr Heynes said. "But we want to make clear that we are on the side of the investor and are helping the drive towards transparency," Mr Heynes added.

certain disclosures is a key ingredient in this analysis. Second, our study contributes to the existing conflict of interest literature by examing a conflict of interest covered by SOX but not covered by the literature: ownership disclosure. Moreover, instead of focusing on whether analyst conflicts of interest lead to skewed recommendations as most of the literature does,[7] we examine the extent to which investors incorporate conflict of interest disclosures in their investment decision. Our study is also unique in this literature because it measures the response to the recommendation and the disclosure in real-time. Third, our paper contributes to the growing literature that examines the types of information individuals use when making their investment decisions (for instance, Huberman and Regev (2001), Barber and Odean (2005), and Engelberg, Sasseville, and Williams (2006)).

We collect data from a unique environment in which disclosures are mandated: television appearances. We choose this setting for two reasons: (1) it is among the few environments in which we can accurately measure the time of the recommendation and the time of the disclosure, and (2) those who follow TV stock-picking advice are likely to be unsophisticated investors and precisely for whom the disclosure requirements are intended.[8] We follow a *Making Money Now* segment of the CNBC lineup, which airs while the U.S. market is open, between March and July of 2005. We record in real time the moment of each analyst recommendation and the moment of each disclosure. Using price, volume and order flow data, we find evidence suggesting that investors do not wait for the disclosure to respond to the recommendation. We find no evidence that investors react when the disclosure information arrives, or that the information revealed during the

[7]One notable exception is Agrawal and Chen (2005) who find that the market recognizes sell-side analyst conflicts and properly discounts optimistic forecasts.

[8]Engelberg, Sasseville, and Williams (2006) provide evidence that those who follow stock-picking advice from the CNBC show *Mad Money* lose money in the short-run to more sophisticated traders.

disclosures has an effect on price, volume or order flow. Our results are robust to various definitions of analyst conflict.

It is important to note that our evidence of no response to the disclosure information does not mean that investors are indifferent as to whether or not ownership disclosures are mandatory. It is possible that what matters to investors is not the information revealed during the disclosures, but the fact that they know that analysts making stock recommendations are required to disclose this information, and that this, in turn, can potentially alter the behavior of analysts. It is thus possible that the mandatory disclosures affect the investment behavior of viewers, regardless of the information disclosed. Because we only have a sample post-SOX, we are unable to rule out this possibility. That is, we are unable to determine whether there is no response to the disclosure information because investors consider these disclosures to be useless, or simply because they do not care about the information content of these disclosures, even though they value the fact that ownership disclosures are mandatory. To make this distinction, we would need to compare the reaction to stock recommendations pre- and post-SOX.

Our data and methodology are most similar to Busse and Green (2002). Busse and Green recorded episodes of the CNBC television shows *Morning Call* and *Midday Call* to determine how quickly the market adjusted to information an analyst first revealed on air. Our study is similar in methodology, but differs significantly in purpose. While the purpose of their paper is to study how quickly market prices react to information in order to examine market efficiency, the purpose of our study is to examine whether the market reacts to the conflict of interest disclosures that follow analyst reports.

The outline of this chapter is as follows: Section 2 describes our data, Section 3 presents our results, and Section 4 concludes.

3.2. Data

CNBC is the major financial network on cable television. Launched on April 17, 1989, CNBC provides news and analysis about financial markets while the market is open and various programming (pre/post market analysis, reruns of NBC shows, talkshows, infomercials, etc.) when the market is closed. Its popularity peaked with the market during late 1999, but it has withstood challenges from other cable financial networks like CNNfn which folded in 2004. Currently, the typical CNBC shows draw anywhere between 100,000 and 300,000 households according to Nielsen Media Research (NMR) but these numbers likely underrepresent CNBC viewership since NMR does not take into account viewers who watch in their office.[9] According to a CNBC spokesman, households that watch CNBC have a median net worth of over $1.3 million.[10]

We collected our data by recording the *Making Money Now* segment of the show *Power Lunch* which airs on CNBC Monday through Friday from 12 p.m. to 2 p.m. EST. Our recordings began on March 29th, 2005, and ended on July 15th, 2005, during which time the S&P 500 rose 5.4%.[11] In the months of April, May, June and July *Power Lunch* drew an average audience of 198,000 households according to monthly reports produced by Nielsen Media Research.[12] As the name suggests, the *Making Money Now* segment

[9]For an example of CNBC ratings in a recent month, see www.mediabistro.com/tvnewser/original/april06ranker.pdf
[10]Source: USA Today, "CNBC's flagship show takes a new tack," December 15, 2005.
[11]During this time period, we observed some instances in which the *Making Money Now* segment did not air due to a variety of news-making instances (e.g. an announcement by President Bush, a hearing with Chairman Greenspan, etc.).
[12]Source: www.mediabistro.com

features money managers who provide stock-picking advice for the viewing audience. Busse and Green (2002) note that money managers may have an incentive to appear on a network like CNBC because the large audience can help build the reputation of the analyst or the firm. Although the *Making Money Now* segment appears in many CNBC programs, its lightning round feature at the end of *Power Lunch* is the most structured. During the lightning round, two money managers make alternating buy and sell recommendations under the time restriction of 15 seconds per recommendation. Although the number of stocks that each money manager actually recommends varies due to the time constraints of the show, each money manager is asked to make 3 buy recommendations and 3 sell recommendations in the lightning round.[13] Following all of the recommendations, the host of the show displays on screen any conflicts of interests the money managers may have: for each stock that is discussed, the screen displays whether or not the money manager personally owns the stock, whether or not the money manager's family owns the stock and whether or not the money manager's firm owns the stock. Also, for sell recommendations, the screen displays whether or not the money manager has a short position. The disclosure data are summarized in Table 3.1. For the buy recommendations, we observe analyst ownership 35.37% of the time, firm ownership 80.85% of the time and no ownership 15.69% of the time.[14] For the sell recommendations, we observe no ownership 78.67% of

[13]For this reason, the number of buy and sell is predetermined to be roughly the same. Therefore, we do not observe the strong disproportion of buy (relative to sell) recommendations documented by others (Elton, Gruber, and Grossman (1986), Stickel (1995) and Malmendier and Shanthikumar (2005)).

[14]These percentages do not sum to 100% since analyst ownership and firm ownership are not mutually exclusive.

the time, a short position 6.05% of the time, analyst ownership 1.44% of the time, and firm ownership 15.27% of the time.[15]

Table 3.1: **Descriptive statistics for the ownership variables.** Reported are the number of instances in which the recommended stock is owned by the analyst's firm, the analyst himself and the analyst's family. We also report the number of short positions and the number of cases in which there was no ownership.

	Buys (376 observations)		Sells (347 observations)	
	Frequency	% of buy observations	Frequency	% of sell observations
Firm	304	80.85	53	15.27
Analyst	133	35.37	5	1.44
Family	18	4.79	0	0
Short Position	0	0	21	6.05
No Ownership	59	15.69	273	78.67

Figure 3.1 illustrates the timing of the recommendation and disclosure events. We recorded the shows with a TIVO machine which has a built-in clock that is set during its routine connections with Network Time Protocol (NTP) servers. Clocks set by NTP servers keep extremely precise time - usually accurate within 10 milliseconds.[16] Using TIVO's clock, we record the time–accurate to the second–of the buy/sell recommendations and ownership disclosures.[17] Over the sample period, we recorded 68 shows with 60 unique analysts who provided a total of 723 recommendations. Table 3.2 gives summary statistics for the recommendations. It indicates that analysts recommended selling and buying firms with similar volatility, size and liquidity but tend to recommend buying

[15]It may seem strange when an analyst's firm owns the security that the analyst recommends selling, but this could be the case if the security is owned in a portfolio of the firm that is not under the analyst's control (for example, if the analyst is a fund manager and his firm owns the security in a fund that the manager does not manage).

[16]See Minar (1999) for an overview of the NTP network and its accuracy.

[17]The recommendations (or disclosures) are typically made on screen, but if they are announced by the analyst or host before they appear on screen, we record this earlier event as the recommendation event (or disclosure event).

Figure 3.1: Illustration of event timing. This figure illustrates the order of events in our data. TV analysts on the show *Making Money Now* make several 15-second buy and sell recommendations. Typically, two analysts on the show each make three buy recommendations and three sell recommendations. Each time an analyst makes a recommendation, we record this as a recommendation event. After all the recommendations have been made, the host of the show discloses whether the analyst, his family, or his firm owns the recommended stocks. Depending on whether the recommended stock was one of the first or one of the last stocks to be recommended, ownership disclosure occurred between 23 seconds and 308 seconds after recommendation.

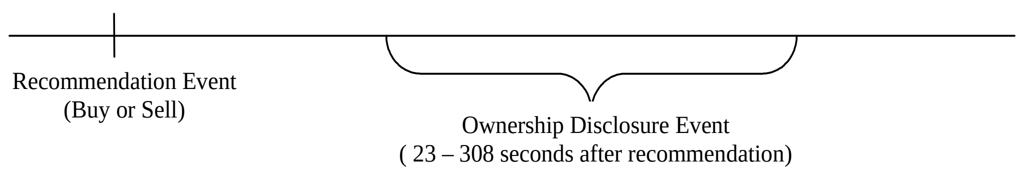

firms with stronger recent performance (the median prior year return is 13% among buys and 5% among sells). Table 3.3 describes the analysts who appeared on the show during our sample period. The table indicates that the vast majority of our analysts are on the buy-side.

Table 3.2: **Recommendation descriptive statistics.** This table reports descriptive statistics on the recommended stocks. Reported are the mean of the log market capitalization on the day before the recommendation (Log Market Cap), the return over the last year (Year Return), the average daily standard deviation over the last year (Year Std Dev), and the average relative spread sampled every minute on the day of the recommendation (Relative Spread). The relative spread is defined as the difference of the bid and ask divided by the midpoint of the bid and ask.

	Buys (376 Observations)			Sells (347 Observations)		
	Mean	Median	Std Dev	Mean	Median	Std Dev
Log Market Cap	9.1	9.2	1.9	9	9.2	1.7
Year Return	0.22	0.13	0.39	0.15	0.05	0.51
Year Std Dev	0.018	0.016	0.009	0.02	0.018	0.009
Relative Spread	8.9	5.5	10.2	8.5	5.5	7.7

Table 3.3: **Description of analysts.**

Name	Company	Title	Buy	Sell	Total
Andrew Seibert	S & T Wealth Management	Senior Portfolio Manager	3	3	6
Barry James	James Advantage Funds	President	9	8	17
Barry Ritholtz	Maxim Group	Chief Market Strategist	7	7	14
Ben Halliburton	Tradition Capital Management	CIO	5	5	10
Ben Pace	Deutsche Bank	CIO of Private Bank Mgmt	1	1	2
Brett Gallagher	Julius Baer Asset Management	Head of Global Equity	6	4	10
Brian Clifford	SunAmerica New Century Fund	Portfolio Manager	9	9	18
Charles Lemonides	Value Works	CIO	14	16	30
Chris Orndorff	Payden and Rygel	Portfolio Manager	3	3	6
Chris Trompeter	Tradition Capital Mgmt	Managing Director	2	2	4
Dan Genter	RNC Genter	CEO	12	11	23
Dan Morgan	Synovus Investment Advisors	Manager	8	7	15
David Dietze	Point View Financial Services	Chief Investment Strategist	3	2	5
David Dreman	Dreman Value Management	Chairman and CIO	3	2	5
David Goerz	Highmark Capital	CIO	6	6	12
David Katz	Matrix Asset Advisors	CIO	3	3	6
David King	Putnam Value fund	Portfolio Manager	11	11	22
Don Hodges	Hodges Fund	President	9	7	16
Doug Altabef	Matrix Asset Advisors	Managing Director	3	2	5
Eric Thorne	Bryn Mawr Trust Wealth Mgmt	Portfolio Manager	14	14	28

Name	Company	Title	Buy	Sell	Total
Gene Henssler	The Henssler financial group	Investment CIO	6	6	12
George Foley	Glenmede Trust	Vice President	9	9	18
Greg Church	Church Capital Mgmt	CIO	2	2	4
Howard Rosencrans	Value Advisory	President	13	14	27
Hugh Johnson	Johnson Illington Advisors	Chairman and CIO	11	10	21
Ivan Feinseth	Matrix USA	Director of Market Research	9	9	18
Jeanne Mockard	George Putnam Fund of Boston	Portfolio Manager	3	3	6
Jeff Kleintop	PNC Advisors	Chief Investment Strategist	3	2	5
Joe Besecker	Emerald Asset Management	President	20	22	42
John Schmitz	Fifth Third Asset Mgmt	Managing Director	2	2	4
Joseph Zock	Capital Mgmt Associates Press	President and Portfolio Mgr	3	3	6
Keith Wirtz	Fifth third asset management	President and CIO	3	3	6
Kevin Divney	Putnam Vista Fund	Co-Portfolio Manager	11	9	20
Lee Schultheis	Alpha Hedge strategies fund	CEO and Chief Invest. Strat.	3	3	6
Malcom Polley	SandT Management	CIO	9	9	18
Mary Lisanti	AH Lisanti Capital Growth	President	3	3	6
Matthew Patsky	Winslow Mgmt	Portfolio Manager	3	2	5
Michael Gallipo	Citizens Funds	Portfolio Manager	3	3	6
Micheal Moe	Think Equity Partners	Chairman and CEO	12	9	21
Mike Blatt	Chemung Canal Trust Company	Senior Portfolio Manager	3	3	6
Nick Colas	Rochdale Research	Director of Research	3	3	6
Noah Blackstein	Dynamic Funds	Portfolio Manager	3	2	5

Name	Company	Title	Buy	Sell	Total
Owen Fitzpatrick	Deutsche Bank	Head of US Equity Group	3	3	6
Paul Noglows	IRG Research	Director of Research	3	3	6
Peter Jankovskis	Oakbrook Investments	Director of Research	3	2	5
Quinn Stills	Palisades Investment Partners	Chairman and CIO	3	0	3
Richard Steinberg	Steinberg Global Asset Mgmt	President and CIO	13	12	25
Rob Lutts	Cabot Money Mgmt	Founder and CIO	3	3	6
Sandy Lincoln	Wayne Hummer Asset Mgmt	CEO	3	2	5
Sarat Sethi	Douglas C. Lane and Associates	Partner and Portfolio Mgr	6	6	12
Scott Rothbort	Lake View Asset Mgmt	President	9	8	17
Shawn Price	Touchstone Large Cap Growth	Co-portfolio Manager	6	6	12
Steve Baeur	Truffle Hound Capital	CEO and CIO	3	3	6
Steve Folker	Fifth Third Asset Mgmt	Managing Director	3	2	5
Steve Neimeth	SunAmerica	Portfolio Manager	8	8	16
Ted Parrish	Henssler Equity Fund	Co-portfolio Manager	3	3	6
Tim Ghriskey	Solaris Group	CIO	11	11	22
Tim Smalls	Execution LLC	Head of US Trading	8	7	15
Tom Laming	Trendstar Advisors	President and CIO	9	8	17
Vince Farrell	Scotsman Capital	Managing Director	9	8	17
	Total # of Analysts: 60		376	347	723

3.3. Analysis

To analyze how trading volume is affected by disclosure, for each recommendation we create the variable $\mathit{VolumeRatio}_{t-1,t}$ to measure how much more trading activity there is in the minutes following the recommendation than in the minutes preceding the recommendation. The variable is defined by the equation,

$$\mathit{VolumeRatio}_{t-1,t} = \frac{Vol_t}{\overline{Vol}},$$

where Vol_t is the volume for the recommended stock between times $t-1$ and t, and $\overline{Vol}$ is the average volume per minute for the recommended stock from $t = -65$ to $t = -5$.[18] Time is measured in minutes, and for now, $t = 0$ represents the time of the recommendation.

VolumeRatio can be used to determine whether investors respond to the recommendation and disclosure events, but it cannot be used to determine whether the trades are buyer or seller induced. To analyze order flow we define the variable $\mathit{Imbalance}_{t-1,t}$ by the equation:

$$\mathit{Imbalance}_{t-1,t} = \frac{Vol_t^B - Vol_t^S}{Vol_t^B + Vol_t^B},$$

where Vol_t^B represents the buyer-initiated volume for the recommended stock between times $t-1$ and t.[19] Vol_t^S is defined analogously for seller-initiated volume.

[18]The recommendation data is matched to quotes and trades data from the New York Stock Exchange Trades and Automated Quotes (TAQ) database.

[19]When constructing our imbalance variable we follow Chordia, Roll, and Subrahmanyam (2001) by using only primary market quotes. Quotes and trades recorded out of sequence or recorded before or after the opening or closing of the market are discarded. We discard quotes with negative bid-ask spreads. We follow Korajczyk and Sadka (2004) and discard observations where the bid-ask spread is above five dollars. Also, we discard observations where the bid-ask spread divided by the midpoint of the quoted bid and

If investors are interested in stock recommendations but wait to hear the disclosures before investing, we would expect post-recommendation volume to be lighter before disclosure occurs and heavier after disclosure occurs. Figure 3.2 plots the mean value of $VolumeRatio_{t-1,t}$ for buy and sell recommendations when the disclosure event is less than 2 minutes ("Early Disclosure") after the recommendation event and when the disclosure event is at least 2 minutes after the recommendation ("Late Disclosure"). Such a partitioning leads to four cases: Buy Recommendation/Early Disclosure, Buy Recommendation/Late Disclosure, Sell Recommendation/Early Disclosure and Sell Recommendation/Late Disclosure.

Figure 3.2 depicts some trading activity immediately after sell recommendations and strong trading activity after buy recommendations, which is consistent with Busse and Green (2002). The figures also provide no evidence that investors wait for disclosure. The figures are quite similar in both the early and late disclosure cases and trading is heaviest in the first minute of the Buy Recommendation/Late Disclosure case when no disclosures have been made.

Figure 3.3 shows the mean cumulative return, the $VolumeRatio_{t-1,t}$ and $Imbalance_{t-1,t}$ when the event is the recommendation and when the event is the disclosure. This figure underscores the fact that price, volume and order-flow all respond to the recommendation event but not the disclosure event.

ask is more than 10% if the midpoint is greater than $50, and we discard observations where the bid-ask spread divided by the midpoint is more than 25% if the midpoint is less than $50. We classify trades as buyer-initiated or seller-initiated trade using the Lee and Ready (1991) algorithm with a one-second lag to match quotes to trades.

Figure 3.2: Volume ratio with early and late disclosure. This figure plots the mean VolumeRatio defined as $\frac{Vol_t}{\overline{Vol}}$ where Vol_t is the volume between times $t-1$ and t and $\overline{Vol}$ is the average volume per minute between times $t=-65$ and $t=-5$. Time $t=0$ corresponds to the recommendation event. The graph is labeled "early disclosure" if disclosures were made less than 2 minutes after the recommendation and "late disclosure" if made after 2 minutes.

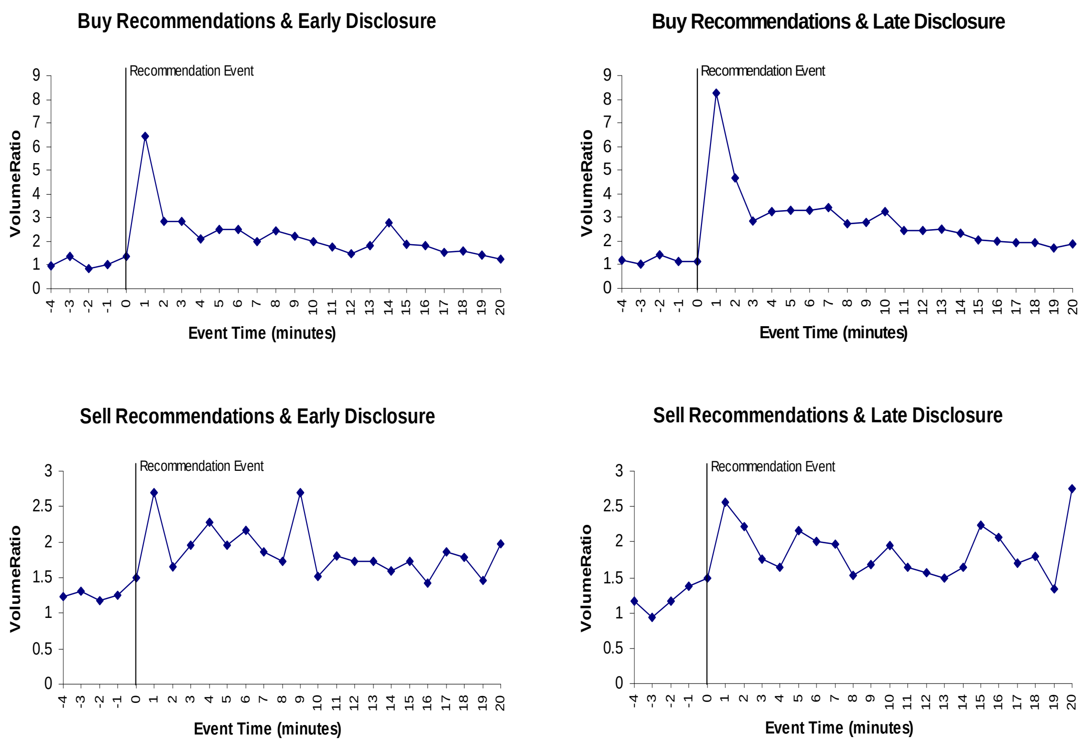

3.3.1. Testable Hypotheses

If investors respond to the disclosure information, we are interested to see how they respond. That is, how do they interpret the information contained in the disclosure. We are interested in testing two hypotheses. The first, which we refer to as the "credibility

Figure 3.3: Cumulative return, volume ratio and imbalance.This figure plots the mean Cumulative Return, Volume Ratio and Imbalance over 25 minutes in event time (5 minutes before the event and 20 after).

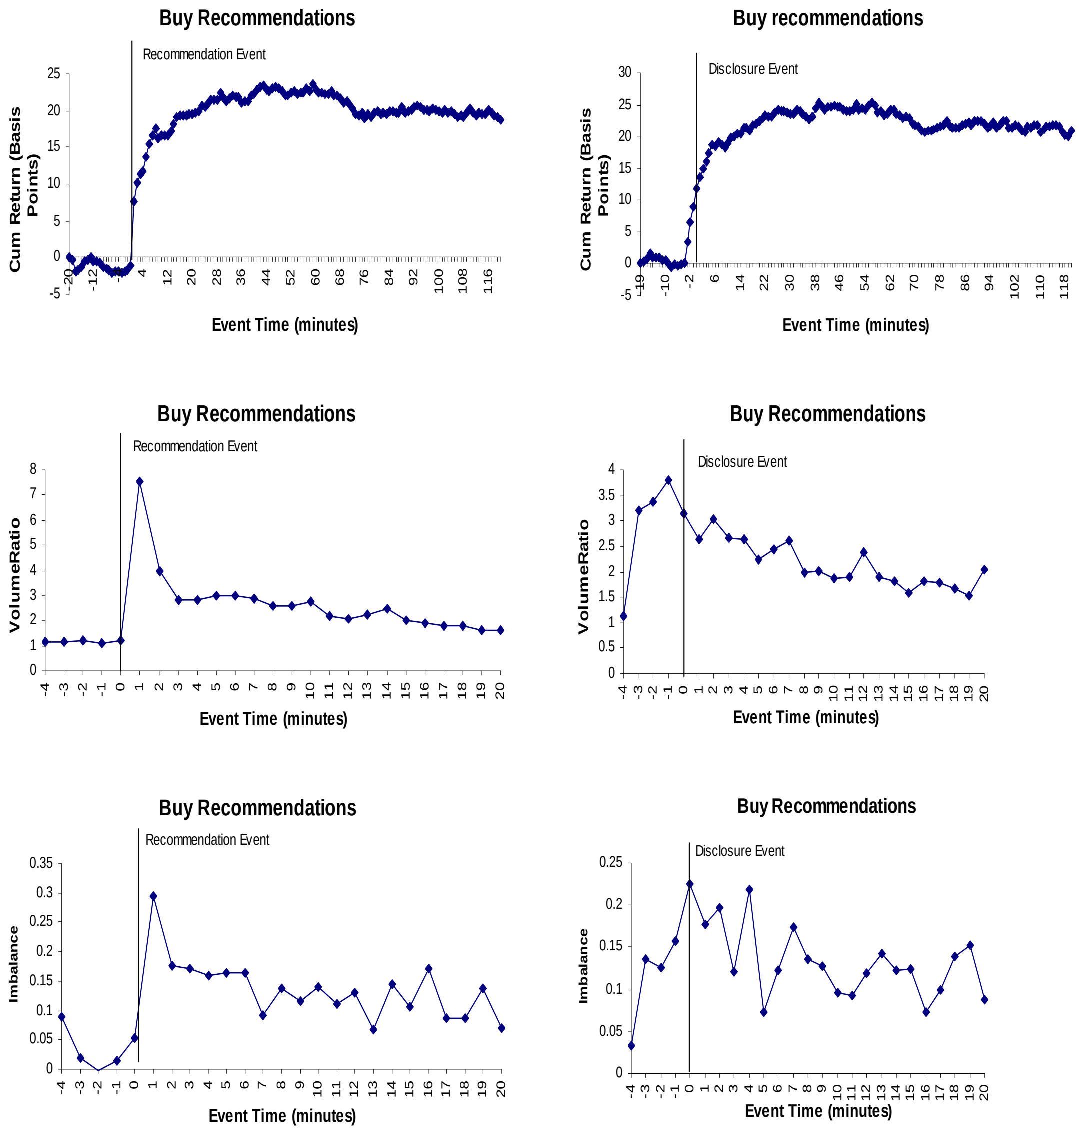

hypothesis", is

$$H_{0,1} \quad : \quad \text{investors consider a stock recommendation to be more credible if the analyst owns the stock.}$$

Under this hypothesis, investors believe that analysts who own the recommended stock genuinely think it is undervalued. In this scenario, the disclosures should be followed by abnormally high buying pressure when the analyst owns the recommended stock. Moreover, this effect should be strongest for the smallest stocks since these are the stocks whose price can be most affected by the analyst's recommendation. Figure 3.4 depicts the credibility hypothesis, after controlling for other variables.

The second, which we refer to as the "vested interest hypothesis", allows for the possibility that investors perceive analysts as acting in their own self-interest. As such, they may discount a buy recommendation made by an analyst who owns the stock because they are concerned that the analyst is attempting to push up the price and therefore increase his personal wealth via the recommendation. The second hypothesis is

$$\begin{aligned} H_{0,2} \quad : \quad & \text{investors discount stock recommendations} \\ & \text{made by an analyst who owns the stock.} \end{aligned}$$

If investors are concerned that analysts are trying to manipulate prices through their recommendations — in particular, if they fear that analysts recommend stocks just to increase the value of their portfolio — there should be less buying pressure following the disclosure when the analyst owns the stock. Moreover, this effect should be strongest for the smallest stocks since these are the stocks whose price can be most affected by the analyst's recommendation. Figure 3.5 depicts the vested interest hypothesis, after controlling for other variables.

We test these hypotheses by regressing returns, buy-sell imbalance, and trading volume on an ownership dummy (either personal or firm ownership), on a size-ownership

Figure 3.4: The credibility hypothesis.

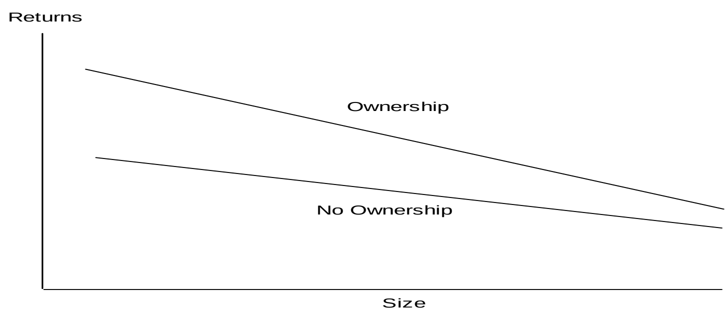

Figure 3.5: The vested interest hypothesis.

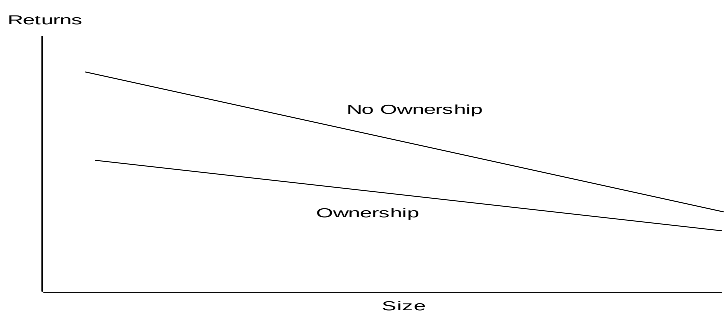

interaction variable, defined as the product of the ownership dummy and the log market capitalization on the day before the recommendation, and on a set of control variables. From Figure 3.4, it is clear that the credibility hypothesis predicts that the coefficient on the ownership dummy variable is positive and that the coefficient on the size-ownership variable is negative in the returns and buy-sell imbalance regressions. From Figure 3.5, it is clear that the vested interest hypothesis predicts that the coefficient on the ownership dummy variables is negative and that the coefficient on the size-ownership variable is positive in the returns and buy-sell imbalance regressions.

3.3.2. Regression Specification

In the following regression analysis, we regress three variables reflecting investors' response to the recommendation — returns, buy-sell imbalance, and trading volume — on an ownership dummy variable (self or firm ownership), a size-ownership interaction variable (defined as the product of the ownership dummy and the log market capitalization of the recommended stock), and a set of control variables. Specifically, we run the regression:

$$y_{i,t} = \alpha + \beta_1 \cdot Ownership_i + \beta_2 \cdot SizeOwnership_i + \beta \cdot x_{i,t} + \varepsilon_{i,t}$$

where $y_{i,t}$ is the return, the buy-sell imbalance or the trading volume of recommended stock i at time t, and $x_{i,t}$ represents a set of control variable, described below. While the credibility hypothesis predicts that $\beta_1 > 0$ and $\beta_2 < 0$, the vested interest hypothesis predicts that $\beta_1 < 0$ and $\beta_2 > 0$.

In our regression analysis, we tried to control for information to which investors may be responding. For most of the recommendations, a chart of the stock's previous one year performance is displayed. To control for this information, we include two of the most prominent features of such charts — the stock's previous one year return and the standard deviation of the stock's one day returns over the past year. In addition, since we expect the response of investors to affect more small and less liquid stocks, we also control for the company's log market capitalization and the relative bid-ask spread.[20,21] For the regressions in disclosure time, we also control for the time between the disclosure and

[20]In regressions not reported, analyst fixed effects were included. The addition of these fixed effects reduces the adjusted R^2, and the F-stats testing whether the coefficients on the fixed effects are jointly zero were insignificant. For this reason, we only include security controls in the following regressions.

[21]Relative spread is defined as the bid-ask spread divided by the average of the bid and the ask.

recommendation events, since the recommendation event has such a significant impact on trading. The regressions are performed in both recommendation time, where $t = 0$ is defined as the time of the recommendation, and in disclosure time, where $t = 0$ is defined as the time of the ownership disclosure. Note that we restrict our attention to buy recommendations because they have significantly larger impacts on price, buy-sell imbalance, and volume than do sell recommendations (see Figure 3.6).

It is not clear how quickly the different types of viewers are able to execute their trades, so, for each of our three variables of interest (returns, volume, and buy-sell imbalance), we look at the variables' values at $t = 1$, $t = 5$, $t = 10$, and $t = 20$, where time is measured in minutes. For each of these times, the "base time" is $t = 0$. That is, when computing the return at $t = 5$, our base price is $t = 0$, and when computing imbalance and volume at $t = 5$, we consider all trades that occur between $t = 0$ and $t = 5$. Notice that the trades executed between $t = 0$ and $t = 1$ are reflected in the variable values at $t = 5$, $t = 10$, and $t = 20$.[22] In this sense, our variables at different times are quite dependent. This allows us to analyze the trades of viewers who are able to quickly execute their trades and the trades of viewers who are slower without splitting these viewers' trades into disjoint time horizons; such splitting would reduce the power of our longer horizon regressions because we would be disregarding the trades of the viewers who are able to quickly execute their trades. We also define returns, volume, and imbalance at $t = 0$. For the $t = 0$ regressions (and only in these regressions), the base time is $t = -5$. We include these regressions to show that there is little evidence that abnormal activity occurs before the recommendation event.

[22]Our use here of cumulative measures of trading activity differs from earlier our analysis where we considered non-cumulative variables to construct Figure 3.2.

Figure 3.6: Cumulative return and spread. This figure plots the cumulative return in recommendation event time where price is defined as (1) the bid, (2) the ask and (3) the midpoint. We also plot the bid-ask spread normalized by price. The dotted lines are the 95% bootstrap confidence intervals. The bootstrap confidence intervals were calculated using identical sample size and 10,000 replications (with replacement) at each time t.

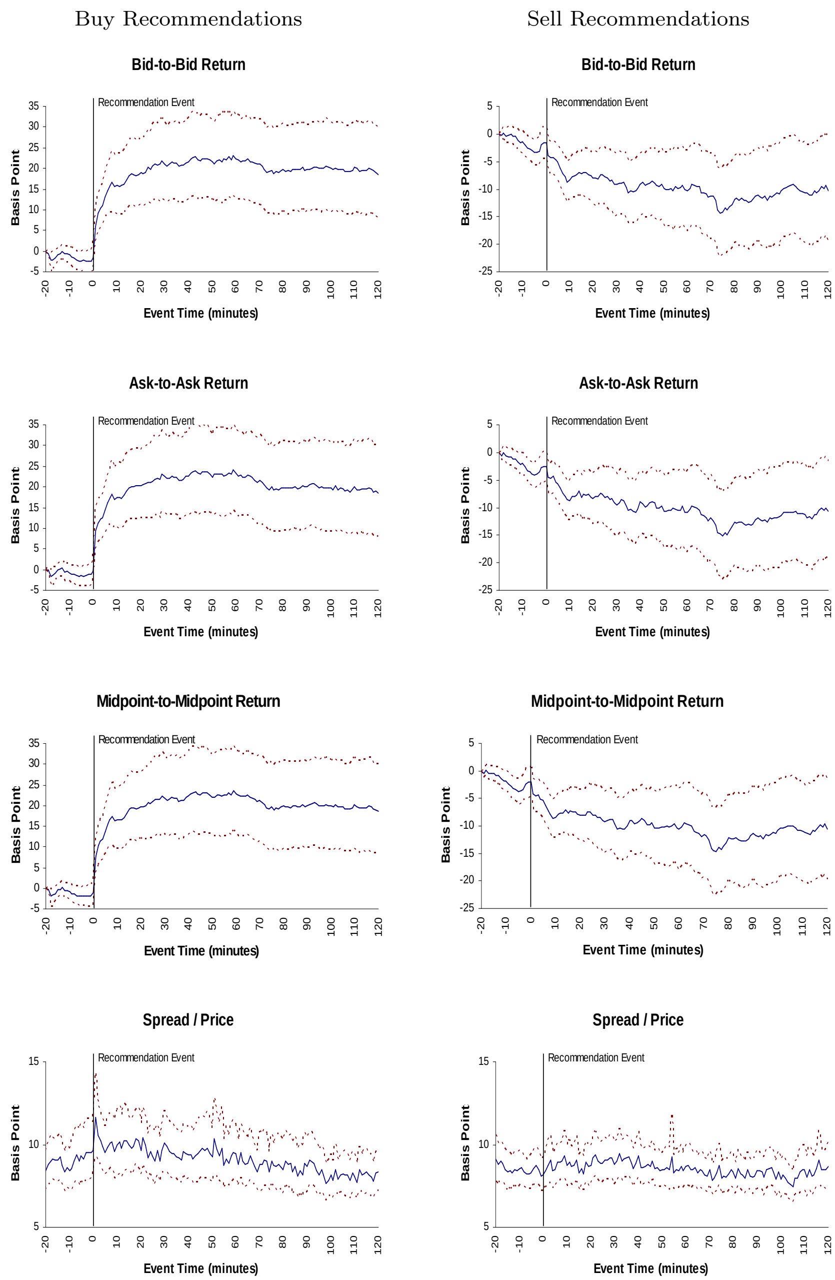

Returns are computed using prices defined as the bid-ask midpoint. Specifically, we have

$$CumulativeReturn_{b,t} = \frac{midpoint_t - midpoint_b}{midpoint_b}$$

where b represents the base time. We use the expression

$$\frac{1}{t-b}\frac{Vol_{b,t}}{\overline{Vol}}$$

as our measure of trading volume, where $Vol_{b,t}$ is the trading volume between the base time and time t, and $\overline{Vol}$ is the average per minute trading volume from $t = -65$ to $t = -5$. Finally, we use

$$Imbalance_{b,t} \quad = \quad \frac{Vol^B_{b,t} - Vol^S_{b,t}}{Vol^B_{b,t} + Vol^S_{b,t}},$$

as our measure of buy-sell imbalance. Vol^B_t represents the buyer-initiated volume for the recommended stock between the base time and time t. Vol^S_t is defined analogously for seller-initiated volume.

3.3.3. Results

The regression results are reported in Tables 3.4 to 3.7. The coefficients and standard errors of the control variables are omitted since our hypothesis only yields predictions for the sign of the coefficients on the ownership dummy and on the size-ownership interaction variable. For each regression we report the adjusted R^2, the robust standard errors, and coefficient estimates of the ownership dummy and size-ownership variable, and the p-value of the F-statistic associated with the null hypothesis that the coefficient on the ownership dummy and size-ownership interaction variable are both equal to zero.

In Table 3.4, we present our regression results using the full sample of buy recommendations. We regress our three variables reflecting investors' response to the recommendation — returns, trading volume, and buy-sell imbalance — on the self ownership dummy variable, on the size-ownership interaction variable, and on our set of control variables.[23]

First consider the regressions in recommendation time. The fact that the R^2 for the volume and returns regressions are significantly higher for the $t = 1$ regressions than for the $t = 0$ ones suggests that we can observe the impact of viewers' trades and that little abnormal activity occurs prior to the recommendation event. The fact that the R^2 actually drops from $t = 0$ to $t = 1$ in the disclosure time regressions suggests that most viewers do not wait for the disclosure event to execute their trades.

Recall that the vested interest hypothesis predicts that the coefficient on the ownership dummy variable is negative and that the coefficient on the size-ownership variable is positive in the returns and buy-sell imbalance regressions at times $t = 1$, $t = 5$, $t = 10$, and $t = 20$. The credibility hypothesis predicts the opposite. To be able to determine which hypothesis holds, it is helpful to first determine whether the coefficients on the ownership variable have the same sign in all of the regressions after $t = 0$. Consider the recommendation time buy-sell imbalance regressions reported in Table 3.4. At $t = 1$ and $t = 5$, the coefficient on the ownership dummy is negative and the coefficient on the size-ownership variable is positive, but at $t = 10$ and $t = 20$ the coefficient on the ownership dummy is positive and the coefficient on the size-ownership variable is negative. This

[23]Recall that our control variables consist of the stock's past one year return, the standard deviation of the stock's one day returns over the past year, the stock's log market capitalization the day before the recommendation, and the stock's relative spread . In the disclosure time regressions we also include the time between the disclosure and recommendation events as a control variable.

switching in sign of the coefficients on the ownership variables makes it difficult to infer in which direction the disclosure information affects investors, if at all.

Since we expect our results to be strongest among the smallest stocks, we run all of our regressions using two sample: the entire sample and the smallest quartile of recommended stocks. The quartiles are based on the market capitalization the day before the recommendation. Table 3.5 shows the results of the regressions of Table 3.4, using the smallest quartile of recommendations. The regressions in Tables 3.6 and 3.7 are identical to the ones in Tables 3.4 and 3.5, except that ownership is defined as firm ownership instead of self ownership. The results presented in Tables 3.5 to 3.7 show that our results are the qualitatively similar whether we use the entire sample or only small stocks, and that the results are robust to a change in the definition of ownership.

Consider Tables 3.4 and 3.5. The coefficients on the self ownership dummy variable and the size-ownership variable are negative and positive, respectively, in the recommendation time returns regressions. This is true when the entire sample is used (Table 3.4) and when the sample is restricted to the smallest quartile of recommendations (Table 3.5). In the disclosure time buy-sell imbalance regressions, however, the signs of the coefficients are reversed: the coefficient on the self ownership dummy variable is positive while the coefficient on the size-ownership variable is negative. Since imbalance and returns both proxy for viewers' desire for the recommended stock, the discrepancy suggests that any significance of the estimates in those regressions is spurious.

Of our 16 sets of returns and buy-sell imbalance regressions (Tables 3.4 to 3.7), we observe "consistency" of the coefficient signs (i.e., the sign of the coefficient on the ownership dummy is the same at $t = 1$, $t = 5$, $t = 10$, and $t = 20$, and the sign of the coefficient on the size-ownership interaction variable is the same at $t = 1$, $t = 5$, $t = 10$, and $t = 20$) five times: four times for self ownership and once for firm ownership. Table 3.8 summarizes the sets of regressions in which we observe consistency. For the regressions where ownership is defined as firm ownership, we observe "switching" in signs of the coefficients for all but one of the regressions in which the dependent variable is returns or imbalance, which suggests that the firm ownership dummy and the size-ownership variable have little effect on returns or buy-sell imbalance.

Overall, our results suggest that the information revealed in the ownership disclosures has little effect on viewers' investment decisions.

3.3.4. Discussion

Our analysis does not support either hypothesis. That is, it does not appear that the "representative" viewer cares whether or not the analyst or his company owns the recommended stock.

The fact that neither hypothesis is supported by the data does not necessarily imply that viewers' trades are not affected by the disclosure information. In reality there are three groups of investors: those who prefer to buy stocks owned by the analyst, those who prefer to buy stocks not owned by the analyst, and those who do not care whether the analyst owns the recommended stock. Our regression analysis suggests that the sizes of the first two groups are similar, and the fact that the disclosure event is not followed by abnormal trading volume suggests that the sizes of the first two groups are small.[24]

[24]Members of the first two group should either (i) wait for the disclosure event before trading, or (ii) unload their positions if the disclosure information was unfavorable from their perspective. In either case, the disclosure event would be followed by heavy trading.

Table 3.4: **Self ownership regressions with size interaction using full sample.** Using our entire sample of recommendations, we regress returns, trading volume, and buy-sell imbalance on a self ownership dummy variable, a self-size interaction variable (defined as the product of the self ownership dummy and the log market cap of the recommended stock) and a set of control variables. Reported are the coefficients and robust standard errors for the ownership dummy and interaction variables, the F-stat for the null that the dummy and interaction are both equal to zero, and the adjusted R^2. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	Disclosure Time					
	Self Ownership		Self-Size Interaction		F p-value	R^2
Return						
t=0	-1.00e-03	[2.2e-03]	1.30e-04	[2.2e-04]	0.26	0.4
t=1	1.40e-03	[1.0e-03]	-1.30e-04	[9.9e-05]	0.27	0.06
t=5	6.70e-04	[1.9e-03]	-5.70e-05	[1.9e-04]	0.83	0.2
t=10	-1.20e-07	[2.2e-03]	1.60e-05	[2.2e-04]	0.81	0.22
t=20	-4.20e-05	[3.7e-03]	2.90e-05	[3.7e-04]	0.76	0.17
Volume						
t=0	0.046	[6.147]	0.033	[0.608]	0.52	0.21
t=1	6.512	[9.870]	-0.589	[0.963]	0.36	0.13
t=5	0.584	[4.611]	-0.023	[0.458]	0.38	0.22
t=10	-0.796	[3.574]	0.100	[0.357]	0.52	0.21
t=20	-0.318	[2.389]	0.047	[0.237]	0.53	0.22
Imbalance						
t=0	-0.433**	[0.213]	0.044**	[0.022]	0.13	0.14
t=1	0.269	[0.399]	-0.026	[0.041]	0.78	0.01
t=5	0.17	[0.281]	-0.018	[0.028]	0.79	0.01
t=10	0.184	[0.204]	-0.022	[0.021]	0.36	0.02
t=20	0.315*	[0.181]	-0.032*	[0.018]	0.22	0.03

	Recommendation Time					
	Self Ownership		Self-Size Interaction		F p-value	R^2
Return						
t=0	-9.30e-05	[6.8e-04]	1.20e-05	[6.8e-05]	0.95	0.02
t=1	-2.30e-03	[1.5e-03]	2.4e-04*	[1.4e-04]	0.06*	0.52
t=5	-5.70e-04	[2.4e-03]	1.00e-04	[2.3e-04]	0.04**	0.41
t=10	-1.00e-03	[3.5e-03]	1.50e-04	[3.5e-04]	0.16	0.41
t=20	-7.40e-04	[3.7e-03]	1.40e-04	[3.6e-04]	0.11	0.43
Volume						
t=0	-0.565	[0.666]	0.036	[0.069]	0.13	0.01
t=1	-4.518	[18.205]	0.594	[1.787]	0.21	0.23
t=5	0.461	[7.755]	0.023	[0.763]	0.21	0.26
t=10	-0.573	[5.450]	0.106	[0.539]	0.21	0.26
t=20	0.089	[3.448]	0.021	[0.342]	0.27	0.25
Imbalance						
t=0	-0.396	[0.366]	0.036	[0.036]	0.47	0.03
t=1	-0.236	[0.265]	0.03	[0.028]	0.43	0.11
t=5	-0.07	[0.210]	0.012	[0.021]	0.41	0.1
t=10	0.062	[0.184]	-0.006	[0.019]	0.94	0.08
t=20	0.205	[0.164]	-0.02	[0.017]	0.45	0.09

Table 3.5: **Self ownership regressions with size interaction using small sample.** Using the smallest quartile of recommendations, we regress returns, trading volume, and buy-sell imbalance on a self ownership dummy variable, a self-size interaction variable (defined as the product of the self ownership dummy and the log market cap of the recommended stock) and a set of control variables. Reported are the coefficients and robust standard errors for the ownership dummy and interaction variables, the F-stat for the null that the dummy and interaction are both equal to zero, and the adjusted R^2. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	Disclosure Time					
	Self Ownership		Self-Size Interaction		F p-value	R^2
Return						
t=0	-3.90e-03	[1.5e-02]	6.20e-04	[2.1e-03]	0.83	0.43
t=1	4.40e-03	[6.2e-03]	-5.50e-04	[8.7e-04]	0.27	0.05
t=5	7.60e-03	[9.8e-03]	-1.10e-03	[1.4e-03]	0.72	0.23
t=10	-6.60e-03	[1.2e-02]	1.10e-03	[1.7e-03]	0.57	0.25
t=20	-1.10e-02	[2.3e-02]	1.80e-03	[3.3e-03]	0.62	0.17
Volume						
t=0	-24.692	[44.648]	3.998	[6.215]	0.10*	0.27
t=1	39.974	[72.257]	-5.014	[9.929]	0.35	0.15
t=5	-20.225	[34.508]	3.253	[4.820]	0.12	0.25
t=10	-17.249	[28.135]	2.64	[3.937]	0.26	0.23
t=20	-15.149	[17.704]	2.335	[2.477]	0.16	0.22
Imbalance						
t=0	0.342	[0.835]	-0.069	[0.128]	0.47	0.06
t=1	1.13	[1.358]	-0.161	[0.199]	0.7	0.04
t=5	1.632	[1.326]	-0.246	[0.193]	0.41	0.06
t=10	0.439	[0.986]	-0.056	[0.143]	0.78	0.02
t=20	0.428	[0.914]	-0.043	[0.131]	0.27	0.04

	Recommendation Time					
	Self Ownership		Self-Size Interaction		F p-value	R^2
Return						
t=0	-2.80e-03	[2.5e-03]	4.40e-04	[3.7e-04]	0.44	0.06
t=1	-1.7e-02*	[8.9e-03]	2.4e-03*	[1.3e-03]	0.17	0.56
t=5	-7.80e-03	[1.4e-02]	1.30e-03	[1.9e-03]	0.42	0.4
t=10	-5.50e-03	[2.1e-02]	9.10e-04	[3.1e-03]	0.79	0.39
t=20	-4.90e-03	[1.6e-02]	8.70e-04	[2.3e-03]	0.71	0.42
Volume						
t=0	-3.528	[2.476]	0.494	[0.347]	0.37	0.04
t=1	-21.549	[84.934]	4.152	[11.699]	0.3	0.23
t=5	-0.853	[38.485]	0.731	[5.395]	0.24	0.29
t=10	-3.8	[28.155]	0.929	[3.988]	0.3	0.28
t=20	-4.408	[18.021]	0.939	[2.562]	0.17	0.26
Imbalance						
t=0	0.126	[2.544]	-0.054	[0.367]	0.35	0.07
t=1	0.991	[0.777]	-0.15	[0.120]	0.44	0.05
t=5	0.956	[0.864]	-0.146	[0.128]	0.51	0.04
t=10	0.384	[0.854]	-0.057	[0.126]	0.9	0.01
t=20	0.008	[0.812]	0.015	[0.119]	0.37	0.05

Table 3.6: **Firm ownership regressions with size interaction using full sample.** Using our entire sample of recommendations, we regress returns, trading volume, and buy-sell imbalance on a firm ownership dummy variable, a self-size interaction variable (defined as the product of the self ownership dummy and the log market cap of the recommended stock) and a set of control variables. Reported are the coefficients and robust standard errors for the ownership dummy and interaction variables, the F-stat for the null that the dummy and interaction are both equal to zero, and the adjusted R^2. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	Disclosure Time					
	Self Ownership		Self-Size Interaction		F p-value	R^2
Return						
t=0	-4.80e-04	[3.3e-03]	7.10e-05	[3.3e-04]	0.75	0.4
t=1	2.40e-03	[1.6e-03]	-2.40e-04	[1.6e-04]	0.32	0.08
t=5	1.00e-03	[4.1e-03]	-5.70e-05	[4.1e-04]	0.38	0.2
t=10	-8.10e-04	[4.8e-03]	1.50e-04	[4.8e-04]	0.34	0.22
t=20	-5.10e-03	[8.1e-03]	5.70e-04	[8.3e-04]	0.62	0.17
Volume						
t=0	3.22	[8.533]	-0.283	[0.849]	0.65	0.21
t=1	6.802	[9.259]	-0.68	[0.911]	0.74	0.13
t=5	6.039	[6.305]	-0.551	[0.624]	0.17	0.23
t=10	5.637	[5.253]	-0.526	[0.521]	0.2	0.22
t=20	2.247	[3.462]	-0.21	[0.343]	0.66	0.23
Imbalance						
t=0	-0.420*	[0.215]	0.045*	[0.024]	0.15	0.14
t=1	0.146	[0.397]	-0.017	[0.043]	0.92	0.01
t=5	0.071	[0.255]	-0.019	[0.028]	0.13	0.02
t=10	0.023	[0.191]	-0.015	[0.020]	0.02**	0.03
t=20	0.198	[0.159]	-0.024	[0.017]	0.3	0.02

	Recommendation Time					
	Self Ownership		Self-Size Interaction		F p-value	R^2
Return						
t=0	-1.10e-03	[8.9e-04]	1.10e-04	[9.0e-05]	0.45	0.03
t=1	-1.10e-03	[2.3e-03]	1.50e-04	[2.3e-04]	0.24	0.52
t=5	-7.50e-04	[4.3e-03]	1.30e-04	[4.3e-04]	0.35	0.41
t=10	2.30e-03	[6.0e-03]	-1.50e-04	[6.0e-04]	0.25	0.41
t=20	-3.70e-03	[6.6e-03]	4.40e-04	[6.6e-04]	0.44	0.43
Volume						
t=0	-1.013	[1.211]	0.039	[0.128]	0.01***	0.05
t=1	-15.745	[26.907]	1.655	[2.691]	0.72	0.23
t=5	-6.74	[10.688]	0.723	[1.079]	0.65	0.26
t=10	-3.379	[8.165]	0.387	[0.824]	0.55	0.26
t=20	-2.237	[5.429]	0.243	[0.549]	0.78	0.25
Imbalance						
t=0	-0.671*	[0.392]	0.077*	[0.041]	0.15	0.04
t=1	-0.006	[0.269]	-0.017	[0.030]	0.05**	0.12
t=5	0.08	[0.214]	-0.012	[0.023]	0.67	0.1
t=10	0.04	[0.158]	-0.013	[0.017]	0.09*	0.09
t=20	0.074	[0.142]	-0.011	[0.015]	0.62	0.09

Table 3.7: **Firm ownership regressions with size interaction using small sample.** Using the smallest quartile of recommendations, we regress returns, trading volume, and buy-sell imbalance on a firm ownership dummy variable, a self-size interaction variable (defined as the product of the self ownership dummy and the log market cap of the recommended stock) and a set of control variables. Reported are the coefficients and robust standard errors for the ownership dummy and interaction variables, the F-stat for the null that the dummy and interaction are both equal to zero, and the adjusted R^2. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	Disclosure Time					
	Self Ownership		Self-Size Interaction		F p-value	R^2
Return						
t=0	7.90e-03	[1.5e-02]	-1.10e-03	[2.2e-03]	0.86	0.43
t=1	1.3e-02*	[7.5e-03]	-1.9e-03*	[1.1e-03]	0.19	0.14
t=5	2.20e-03	[1.7e-02]	-1.80e-04	[2.4e-03]	0.7	0.24
t=10	-1.40e-03	[1.9e-02]	3.30e-04	[2.7e-03]	0.81	0.24
t=20	-3.70e-02	[4.4e-02]	5.60e-03	[6.4e-03]	0.59	0.19
Volume						
t=0	35.436	[51.486]	-4.716	[7.153]	0.61	0.28
t=1	86.254	[55.459]	-12.175	[7.742]	0.29	0.19
t=5	58.861	[39.650]	-8.162	[5.532]	0.33	0.3
t=10	47.603	[33.621]	-6.576	[4.678]	0.36	0.29
t=20	24.392	[21.801]	-3.393	[3.038]	0.54	0.25
Imbalance						
t=0	0.759	[0.712]	-0.128	[0.108]	0.32	0.06
t=1	1.116	[1.346]	-0.141	[0.208]	0.34	0.05
t=5	-0.274	[0.979]	0.033	[0.144]	0.84	0.04
t=10	0.048	[0.753]	-0.011	[0.112]	0.95	0.02
t=20	0.792	[0.733]	-0.105	[0.110]	0.24	0.04

	Recommendation Time					
	Self Ownership		Self-Size Interaction		F p-value	R^2
Return						
t=0	1.40e-03	[3.6e-03]	-3.30e-04	[5.5e-04]	0.2	0.09
t=1	9.70e-04	[1.1e-02]	-1.10e-04	[1.5e-03]	0.94	0.55
t=5	1.70e-03	[1.9e-02]	-1.20e-04	[2.7e-03]	0.82	0.4
t=10	1.70e-02	[2.4e-02]	-2.20e-03	[3.5e-03]	0.57	0.4
t=20	-6.60e-03	[2.8e-02]	1.10e-03	[4.0e-03]	0.91	0.42
Volume						
t=0	-4.592	[4.195]	0.559	[0.591]	0.18	0.11
t=1	-11.463	[111.198]	2.122	[15.402]	0.85	0.23
t=5	-5.971	[45.351]	1.097	[6.314]	0.81	0.28
t=10	0.066	[34.968]	0.238	[4.837]	0.7	0.27
t=20	-1.381	[24.077]	0.348	[3.346]	0.75	0.25
Imbalance						
t=0	-2.567	[1.622]	0.354	[0.242]	0.17	0.09
t=1	0.955	[0.711]	-0.155	[0.106]	0.28	0.06
t=5	0.358	[0.815]	-0.046	[0.122]	0.77	0.03
t=10	-0.222	[0.629]	0.034	[0.094]	0.94	0.01
t=20	0.7	[0.614]	-0.098	[0.093]	0.39	0.05

Table 3.8: **Summary of the regression results presented in Tables 3.4 to 3.7.** Here, "consistency" occurs if the coefficients of the ownership dummy and interaction variables have the same sign at $t = 1, 5, 10$ and 20. The "Variable Signs" column presents the sign of the dummy and interaction variables (in that order) when the coefficients are consistent across time. When the coefficients are not consistent, "-" is reported. For example, "$(-, +)$" is reported for the self ownership regressions with the return dependent variable using the full sample in disclosure time because the coefficient for the dummy variable is negative and the coefficient for the interaction variable is positive in these regressions at $t = 1, 5, 10$ and 20.

	Self Ownership Regressions			
	Sample	Event Time	Consistency	Variable Signs
Return	full	disclosure	no	-
	full	recommendation	yes	(-,+)
	small	disclosure	no	-
	small	recommendation	yes	(-,+)
Imbalance	full	disclosure	yes	(+,-)
	full	recommendation	no	-
	small	disclosure	yes	(+,-)
	small	recommendation	no	-

	Firm Ownership Regressions			
	Sample	Event Time	Consistency	Variable Signs
Return	full	disclosure	no	-
	full	recommendation	no	-
	small	disclosure	no	-
	small	recommendation	no	-
Imbalance	full	disclosure	yes	(+,-)
	full	recommendation	no	-
	small	disclosure	no	-
	small	recommendation	no	-

Hence, our results suggest that investors are not influenced by the information contained in the ownership disclosures.[25]

It is important to stress that our evidence of no response to the disclosure information does not mean that investors are indifferent as to whether or not ownership disclosures are mandatory. It is possible that what matters to investors is not the information revealed during the disclosures, but the fact that analysts making stock recommendations are required to disclose this information. It is thus possible that the mandatory disclosures give investors greater confidence in analysts' recommendations, regardless of the information disclosed. Because we only have a sample post-SOX, we are unable to rule out this possibility. That is, we are unable to determine whether there is no response to the disclosure information because investors consider these disclosures to be useless, or simply because they do not care about the information content of these disclosures, even though they value the fact that ownership disclosures are mandatory. To make this distinction, we would need to compare the reaction to stock recommendations pre- and post-SOX.

3.4. Conclusion

We analyzed returns, buy-sell imbalance, and trading volume immediately following stock recommendations and ownership disclosures made on a *Making Money Now* segment of the *CNBC* show *Power Lunch*. We find a significant increase in trading volume, returns, and buy-sell imbalance immediately following stock recommendations, but no such dramatic change following ownership disclosures. Furthermore, we find no evidence

[25]We are aware of the possibility that the disclosure event affects investors' behavior but that we cannot detect it because our tests are not powerful enough. In general, it is impossible to disprove a prediction that the means of two random variables are unequal. Hence, our results are merely suggestive.

that investors respond to the information contained in the ownership disclosures. Unfortunately, we are unable to determine whether there is no response to the disclosure information because investors consider these disclosures to be useless, or simply because they do not care about the information content of these disclosures, even though they value the fact that ownership disclosures are mandatory. Our results could, nonetheless, be useful to policymakers who must balance the costs imposed on those who must disclose with the benefits to individual investors for whom the disclosures are intended.

References

Abel, Andrew B., and Janice C. Eberly, 1994, A unified model of investment under uncertainty, *American Economic Review* 84, 1369–1384.

——— , 2002, Investment and q with fixed costs: An empirical analysis, *Working paper.*

Agrawal, A., and M. Chen, 2005, Do analyst conflicts matter: Evidence from stock recommendations, *Working paper, University of Alabama and University of Maryland.*

Barber, B., and T. Odean, 2005, All that glitters: The effect of attention on news on the buying behavior of individual and institutional investors, *Working paper, University of California, Davis* pp. Working paper, University of California, Davis.

Barnett, Steven, and Plutarchos Sakellaris, 1998, Non-linear response of firm investment to q: Testing a model of convex and non-convex adjustment costs, *Journal of Monetary Economics* 42, 261–288.

Brennan, Michael J., 1958, The supply of storage, *American Economic Review* 48, 50–72.

——— , 1990, Latent assets, *Journal of Finance* 45, 709–730.

——— , 1991, The price of convenience and the valuation of commodity contingent claims, in D. Lund, and B. Oskendal, ed.: *Stochastic Models and Option Values* . pp. 33–71 (Elsevier Science Publishers B.V., North-Holland).

Busse, J., and C. Green, 2002, Market efficiency in real time, *Journal of financial economics* 65, 415–437.

Carlson, Murray, Zeigham Khokher, and Sheridan Titman, 2002, An equilibrium analysis of exhaustible resource investments, *Working Paper.*

Casassus, Jaime, and Pierre Collin-Dufresne, 2005, Stochastic convenience yield implied from commodity futures and interest rates, *Journal of Finance* 60, 2283–2331.

——— , and Bryan R. Routledge, 2004, Equilibrium commodity prices with irreversible investment and non-linear technologies, *Working Paper*.

Chambers, Marcus J., and Roy E. Bailey, 1996, A theory of commodity price fluctuations, *Journal of Political Economy* 104, 924–957.

Chen, Long, Pierre Collin-Dufresne, and Robert Goldstein, 2004, On the relation between credit spread puzzles and the equity premium puzzle, *Working Paper*.

Chordia, T., R. Roll, and A. Subrahmanyam, 2001, Market liquidity and trading activity, *Journal of finance* 56, 501–530.

Cybenko, G., 1989, Approximation by superpositions of a sigmoidal function, *Mathematical control signals systems* 2, 303–314.

Deaton, Angus, and Guy Laroque, 1992, On the behaviour of commodity prices, *Review of Economic Studies* 59, 1–23.

——— , 1996, Competitive storage and commodity price dynamics, *Journal of Political Economy* 104, 896–923.

Doms, Mark, and Timothy Dunne, 1998, Capital adjustment patterns in manufacturing plants, *Review of economic dynamics* 1, 409–429.

Dugar, A., and S. Nathan, 1995, The effect of investment banking relationships on financial analysts' earnings forecasts and investment recommendations, *Contemporary Accounting Research* 12, 131–160.

Elton, E., M. Gruber, and S. Grossman, 1986, Discrete expectational data and portfolio performance, *Journal of Finance* 46, 699–714.

Engelberg, J., C. Sasseville, and J. Williams, 2006, Is the market mad: Evidence from mad money, *Working paper, Northwestern University*.

Fama, Eugene F., and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–25.

Funahashi, K.-I., 1989, On the approximate realization of identity mappings by neural networks, *Neural Networks* 2, 183–192.

Gibson, Rajna, and Eduardo S. Schwartz, 1990, Stochastic convenience yield and the pricing of oil contingent claims, *Journal of Finance* 45, 959–976.

Hayashi, Fumio, 1982, Tobin's marginal and average q: A neoclassical interpretation, *Econometrica* 50, 213–224.

Hornik, K., M. Stinchcombe, and H. White, 1989, Multilayer feedforward networks are universal approximators, *Neural Networks* 2, 359–366.

Huberman, G., and T. Regev, 2001, Contagious speculation and a cure for cancer: A nonevent that made stock prices soar, *Journal of Finance* 56, 387–396.

Judd, Kenneth L., 1998, *Numerical Methods in Economics* (The MIT Press).

Kaldor, Nicholas, 1939, Speculation and economic stability, *The Review of Economic Studies* 7, 1–27.

Kogan, Leonid, Dimitry Livdan, and Amir Yaron, 2005, Futures prices in a production economy with investment constraints, *Working Paper*.

Korajczyk, R., and R. Sadka, 2004, Are momentum profits robust to trading costs, *Journal of Finance* 59, 1039–1082.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2006, The convenience yield on treasury debt, *Working Paper*.

Lee, C., and M. Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance* 45, 733–746.

Lin, H., and M. McNichols, 1998, Underwriting relationships, analysts' earnings forecasts, and investment recommendations, *Journal of Accounting and Economics* 25, 101–127.

Malmendier, U., and D. Shanthikumar, 2005, Are investors naive about incentives?, *Working paper, Stanford University*.

Michaely, R., and K. Womack, 1999, Conflict of interest and the credibility of underwriter analyst recommendations, *Review of Financial Studies* 12, 653–686.

Minar, N., 1999, A survey of NTP network, *MIT Media Lab*.

Miranda, Mario J., and Paul L. Fackler, 2002, *Applied Computational Economics and Finance* (The MIT Press).

Mussa, Micheal, 1977, External and internal adjustment costs and the theory of aggregate and firm investment, *Economica* 44, 163–178.

Niederreiter, Harald, 1988, Low-discrepancy and low-dispersion sequences, *Journal of Number Theory* 30, 51–70.

Reed, Russell D., and Robert J. Marks, 1999, *Neural Smithing: Supervised Learning in Feedforward Artificial Neural Networks* (The MIT Press).

Routledge, Bryan R., Duane J. Seppi, and Chester S. Spatt, 2000, Equilibrium forward curves for commodities, *Journal of Finance* 55, 1297–1337.

Scheinkman, Jose A., and Jack Schechtman, 1983, A simple competitive model with production and storage, *Review of Economic Studies* 50, 427–441.

Schwartz, Eduardo, 1997, The stochastic behavior of commodity prices: Implications for valuation and hedging, *Journal of Finance* 52, 923–973.

——— , and James E. Smith, 2000, Short-term variations and long-term dynamics in commodity prices, *Management Science* 46, 893–911.

Stickel, S., 1995, The anatomy of the performance of buy and sell recommendations, *Financial Analysts Journal* 51, 25–39.

Telser, Lester G., 1958, Futures trading and the theory of storage of cotton and wheat, *Journal of Political Economy* 66, 233–255.

Tobin, James, 1969, A general equilibrium approach to monetary theory, *Journal of Money, Credit and Banking* 1, 15–29.

Working, Holbrook, 1949, The theory of the price of storage, *American Economic Review* 39, 1254–1262.

APPENDIX A

Numerical Procedure

Traditionally, dynamic programming problems using recursions are solved by discretizing the state space and then solving the Bellman equation at each point on the grid. For the value one step ahead, the value at the grid point closest to the value of the state variables one period ahead is used. However, such an approach is subject to the curse of dimensionality since the size of the grid used to solve a problem and have reasonable degree of accuracy can grow exponentially with the dimensionality of the problem. Since our problem involves six state variables, we decided to turn to functional approximation.

While polynomial approximation schemes are more common function approximators[1], we opted for a different method since these methods are highly subject to the curse of dimensionality. Guiding our choice primarily by a consideration of the computational requirements (time and memory) of a given method in higher dimensions, we chose to approximate the value function using feedforward neural networks. Cybenko (1989), Hornik, Stinchcombe, and White (1989), and Funahashi (1989) prove that neural networks with one-hidden layer can approximate any continuous function to any level of precision if the transfer function is of sigmoidal type and the transfer function in the output layer is linear. Although, these theorems do not say anything about which weights and biases

[1]Judd (1998) and Miranda and Fackler (2002) provide an excellent overview of function approximation methods.

to use, or how many nodes to use in the hidden layer, they do say something about the potential of neural networks to approximate functions to the desired degree of accuracy.

Because neural networks try to capture the underlying systematic aspects of the data rather than to fit the specific points, which is sometimes equivalent to fitting the noise in the data, neural networks are powerful functional approximators. However, overfitting is still possible if too many training cycles are used. Since it is not possible a priori to determine the optimal number of training cycles, we address this issue by having two samples: a training sample and a testing sample. The parameters are estimated using the training sample and the quality of the fit is measured using the testing sample. That is, once we obtain the parameter estimates from the training sample, we compute the approximate function values using the inputs from the testing sample. We then compute the error between these estimated values and the target values. At each time step, we obtain neural network estimates for different numbers of nodes and choose the estimate which has the lowest error in the testing sample. Also, since the error function to be minimized when training the neural network often has many local minima (Judd, 1998), we give different starting values when training the neural network and compare the fit at different staring values. Finally, initializing the staring value to small random numbers may help with the training of the neural network if sigmoidal saturation is likely (see Reed and Marks (1999) for a discussion).

A.1. Functional approximation using neural networks

We use a neural network with one hidden layer, using the tanh transfer function in the hidden layer and a linear transfer function in the output layer.[2] Let d be the number of state variables (in this case, $d = 6$). Let N be the number of grid points. Let w be the number of nodes in the hidden layer. The neural network is trained to find the parameters $\alpha\ (w \times 1)\,, B\ (w \times d)\,, \theta\ (w \times 1)$ and $\gamma\ (1 \times 1)$ that best approximate (in the mse sense) the target function (say, $f(x)$):

$$f(x) \approx \hat{f}(x) = \alpha' \tanh(Bx + \theta) + \gamma$$

where x is a $d \times N$ vector containing the state variables. The number of nodes used at each time step is determined by the best fit in the testing sample.

A.2. Steps in value function iteration

The value function iteration can best be described in six steps:

(1) Choose a stopping criterion η.

(2) Start with an initial guess for the value function, V^i with $i = 0$.

(3) Use a neural network to approximate the V^i. That is, find α^i , B^i , θ^i , and γ^i such that

$$V^i \approx \hat{V}^i(x) = \alpha^{i\prime} \tanh\left(B^i x + \theta^i\right) + \gamma^i$$

[2]Our results were not affected by whether we approximate the value function using one or two hidden layers. For this reason, we did not see the need to have an additional hidden layer.

where x is a $6 \times N$ matrix containing the state variables at all of the grid points

$$\left(\text{i.e. } x = \begin{bmatrix} z_1 & z_2 & \ldots & z_N \\ \varepsilon_1 & \varepsilon_2 & \ldots & \varepsilon_N \\ \varsigma_1 & \varsigma_2 & \ldots & \varsigma_N \\ Q_1 & Q_2 & \ldots & Q_N \\ K_1 & K_2 & \ldots & K_N \\ W_1 & W_2 & \ldots & W_N \end{bmatrix} \right)$$

.

(4) Solve the agent's maximization problems, using $\hat{V}^i$ to compute the continuation value, to obtain V^{i+1}. That is,

$$V^{i+1}(x) = \max_{S,I,C} u(C, \varsigma) + \beta E_t \left[\hat{V}^i(x_{t+1}) \right]$$

where x_{t+1} is the matrix containing the value of the state variables one period ahead and $E_t \left[\hat{V}^i(x_{t+1}) \right]$ is the expectation of the value function one period ahead.

(5) Use a neural network to approximate V^{i+1}. That is, find α^{i+1} , B^{i+1} , θ^{i+1}, and γ^{i+1} such that

$$V^{i+1} \approx \hat{V}^{i+1}(x) = \alpha^{i+1\prime} \tanh\left(B^{i+1} x + \theta^{i+1}\right) + \gamma^{i+1}$$

where x is a $6 \times N$ matrix containing the state variables at all of the grid points.

(6) Continue iterating until $\left\| \hat{V}^{i+1} - V^i \right\| < \eta$.

APPENDIX B

Prepaid Forward Price and Convenience Yield

Consider the agent's optimization problem:

$$V(\Lambda_t) = \max_{S_t, C_t} u(C_t) + \beta E_t [V(\Lambda_{t+1})]$$

s.t.

$$\begin{aligned} Q_{t+1} &= Q_t + f(\varepsilon_t, K_t) - S_t \\ K_{t+1} &= K_t(1-\delta) \\ W_{t+1} &= W_t + p_t S_t - C_t \\ S_t &\leq Q_t + f(K_t, \varepsilon_t) \\ C_t &\leq W_t + p_t S_t \end{aligned}$$

where $\Lambda_t = (z_t, \varepsilon_t, \varsigma_t, Q_t, K_t, W_t)$. The Lagrangian for this problem is

$$Ł = \max_{S_t, C_t} u(C_t) + \beta E_t [V(\Lambda_{t+1})] - \lambda_t (S_t - Q_t - f(K_t, \varepsilon_t)) - \mu_t (C_t - W_t - p_t S_t)$$

s.t.

$$
\begin{aligned}
Q_{t+1} &= Q_t + f(\varepsilon_t, K_t) - S_t \\
K_{t+1} &= K_t(1-\delta) \\
W_{t+1} &= W_t + p_t S_t - C_t \\
\lambda_t, \mu_t &\geq 0
\end{aligned}
$$

where λ_t and μ_t are Lagrangian multipliers. This gives the two first-order conditions

$$p_t\left(\beta E_t\left[V_W\left(\Lambda_{t+1}\right)\right] + \mu_t\right) = \beta E_t\left[V_Q\left(\Lambda_{t+1}\right)\right] + \lambda_t \tag{B.1}$$

$$u'\left(C_t\right) = \beta E_t\left[V_W\left(\Lambda_{t+1}\right)\right] + \mu_t \tag{B.2}$$

and the Kuhn-Tucker conditions

$$
\begin{aligned}
\lambda_t\left(S_t - Q_t - f\left(K_t, \varepsilon_t\right)\right) &= 0 \\
\mu_t\left(C_t - W_t - p_t S_t\right) &= 0.
\end{aligned}
$$

We also have

$$V_W\left(\Lambda_t\right) = \beta E_t\left[V_W\left(\Lambda_{t+1}\right)\right] + \mu_t \tag{B.3}$$

$$V_Q\left(\Lambda_t\right) = \beta E_t\left[V_Q\left(\Lambda_{t+1}\right)\right] + \lambda_t. \tag{B.4}$$

Combining equations (B.2) and (B.3), we have that

$$u'(C_t) = \beta^{n+1} E_t \left[V_W \left(\Lambda_{t+n+1} \right) \right] + \sum_{i=0}^{n} \beta^i E_t \left[\mu_{t+i} \right]. \tag{B.5}$$

Using equation (B.5), we can write the n-period risk-free rate, $r_{t,t+n}$, as a function of the Lagrangian multipliers for the liquidity constraint from t up to $t+n$:

$$\begin{aligned}
\frac{1}{\left(1 + r_{t,t+n}\right)^n} &= E_t \left[\beta^n \frac{u'\left(C_{t+n}\right)}{u'\left(C_t\right)} \right] \\
&= E_t \left[\frac{\beta^n \left(\beta E_{t+n} \left[V_W \left(\Lambda_{t+n+1} \right) \right] + E_{t+n} \left[\mu_{t+n} \right] \right)}{\beta^{n+1} E_t \left[V_W \left(\Lambda_{t+n+1} \right) \right] + \sum_{i=0}^{n} \beta^i E_t \left[\mu_{t+i} \right]} \right] \\
&= \frac{\beta^{n+1} E_t \left[V_W \left(\Lambda_{t+n+1} \right) + \beta^n E_t \left[\mu_{t+n} \right] \right]}{\beta^{n+1} E_t \left[V_W \left(\Lambda_{t+n+1} \right) \right] + \sum_{i=0}^{n} \beta^i E_t \left[\mu_{t+i} \right]} \\
&= \frac{1}{1 + \frac{\sum_{i=0}^{n-1} \beta^i E_t\left[\mu_{t+i}\right]}{\beta^n E_t\left[u'\left(C_{t+n}\right)\right]}} \\
\Rightarrow \quad r_{t,t+n} &= \left(1 + \frac{\sum_{i=0}^{n-1} \beta^i E_t \left[\mu_{t+i} \right]}{\beta^n E_t \left[u'\left(C_{t+n}\right) \right]} \right)^{\frac{1}{n}} - 1.
\end{aligned}$$

Combining equations (B.1) and (B.2), we have

$$p_t = E_t \left[\beta \frac{u'\left(C_{t+1}\right)}{u'\left(C_t\right)} p_{t+1} \right] + \frac{\lambda_t}{u'\left(C_t\right)}. \tag{B.6}$$

Using equations (B.4), (B.1), and (B.5), we can obtain an expression for the n-period shadow prepaid forward price, $F^P_{t,t+n}$:

$$\begin{aligned} F^P_{t,t+n} &= \frac{\beta^n E_t\left[V_Q\left(\Lambda_{t+n}\right)\right]}{u'\left(C_t\right)} \\ &= \frac{\beta^n E_t\left[\beta E_{t+n}\left[V_Q\left(\Lambda_{t+n+1}\right)\right]+\lambda_{t+n}\right]}{u'\left(C_t\right)} \\ &= \frac{\beta^n E_t\left[p_{t+n}\left(\beta E_{t+n}\left[V_W\left(\Lambda_{t+n+1}\right)\right]+\mu_{t+n}\right)\right]}{u'\left(C_t\right)} \\ &= \beta^n E_t\left[\frac{u'\left(C_{t+n}\right)}{u'\left(C_t\right)} p_{t+n}\right]. \end{aligned} \tag{B.7}$$

Using equations (B.6) and (B.7), we can express the n-period convenience yield, $q_{t,t+n}$, as a function of the Lagrangian multipliers for the commodity inventory constraint from t up to $t+n$:

$$\begin{aligned} F^P_{t,t+n} &= p_t\left(1-\sum_{i=0}^{n-1}\frac{\beta^i E_t\left[\lambda_{t+i}\right]}{p_t u'\left(C_t\right)}\right) \\ &= \frac{p_t}{\left(1+q_{t,t+n}\right)^n} \\ \Rightarrow & \quad q_{t,t+n} = \left(\frac{p_t u'\left(C_t\right)}{p_t u'\left(C_t\right)-\sum_{i=0}^{n-1}\beta^i E_t\left[\lambda_{t+i}\right]}\right)^{\frac{1}{n}}-1. \end{aligned}$$