Illiquidity Costs In A

Discrete Time Model

by

Ryan Domsy

An essay submitted to the Department of Economics

in partial fulfillment of the requirements for

the degree of Master of Arts

Queen's University

Kingston, Ontario, Canada

December 2008

Acknowledgements

I would like to thank Frank Milne for all he has taught me during my time as a graduate student and for his help and guidance which made this paper possible.

Table of Contents

1. Introduction

Uncertainty in financial markets creates the ability to make large profits while also subjecting investors to significant risks. To cope with the risks inherent in financial markets an investor must make use of sound risk management practices. These practices help investors better understand financial risks and be better equipped to deal with them. Risk management has developed very strong models to deal with certain financial risks. However, the research performed on liquidity risk has yet to lead to a comprehensive solution, leaving room for a great deal of research. This research is very important since liquidity risk can have severe effects on the execution price of a transaction leading to extreme losses.

The traditional models of intertemporal portfolio choice (e.g. Merton (1969), Merton (1971) and Cox and Huang (1989)) make many fundamental assumptions. Three of these are very important in the discussion of liquidity risk; Market completeness, unbounded transaction size and the ability to continuously trade assets. These assumptions are very useful to prove closed form solutions to these models and also provide a foundation for standard option pricing theory (e.g. Black and Scholes (1973) and Merton (1973)). Although these assumptions have proven to be useful they are unrealistic when not restricted to market risk. An interesting result of liquidity risk is that it causes these assumptions to fail.

Liquidity risk can put a constraint on the price a trade can be executed at and can also restrict the volume of the trade. These two situations have proven to become very extreme when certain conditions arise. If a market is becoming illiquid, any transaction (even of minimal size) starts to cause large impacts on an assets price. The result of this is that in some cases assets may only be sold for a fraction of their ‘value’. The other situation is where markets become so illiquid the market evaporates. Here an investor is unable to find a counterparty to take on a trade of any size. In both of these scenarios the assumptions of complete and unrestricted markets have failed due to liquidity risk.

A current example of liquidity risk can be seen in the market for mortgage-backed securities (MBS). The bursting of the bubble in the US housing market, starting in 2007, caused homeowners to walk away from their mortgages. This happened because the market value of their home fell below the remaining value of their mortgage and many homeowners were not able to handle the payments on their sub-prime mortgages. As this happened, the number of mortgages packaged in MBSs that were defaulting increased notably. The increases in defaults made the MBSs become extremely risky (highly volatile) even though investors had previously thought they were safe. As the volatility increased, the price investors were willing to pay for a MBS plummeted and the market dried up and exhibited large bid-ask spreads (Brunnermeier 2008). The illiquidity in this market has caused losses in billions of dollars for financial institutions. In some cases the losses were so severe the institutions were forced to release large percentages of their employees and, in the most extreme cases, the financial institutions were sold at extreme discounts in an attempt to avoid bank failures.

Another important real world display of liquidity costs happened in the late nineties and as a result many hedge funds suffered very large losses. This occurred because they had large speculative positions in various bond markets. These markets collapsed as a result of two forces; the Asian financial crisis (began in mid 1997) and the Russian financial crisis (began in mid 1998). These two forces collapsed bond markets in Russia (government bonds defaulted), Japan and Europe as investors liquidated positions to move funds to the much safer US treasury bonds. Hedge fund traders attempted to liquidate their large speculative positions so they could meet their margin calls as bond prices drastically fell. When these firms, most notably Long-Term Capital Management (Lowenstein (2001)), tried to undo their positions the traders quickly realized there was no market. The market had disappeared overnight resulting in extreme losses being realized by these hedge funds and the eventual failure of certain firms.

In both of the situations discussed markets suffered from severe illiquidity. The illiquidity that dominated the markets caused abnormal losses and even led to the failure of financial institutions. This demonstrates the need of a manager to properly account for the potential costs of illiquidity. As a result of these rare events, liquidity risk has experienced increased attention. This has been very useful because the modeling of this risk is still preliminary. This contrasts with the literature on the other financial risks, market and credit risk, which have been modeled much more extensively. Academics and practitioners have been in search of methods to identify, model and respond to liquidity risk.

Empirical research performed in the early 1990's has shown that assets of very similar nature exhibit variations in price which can be attributed to the assets level of liquidity. Amihud and Mendelson (1991) showed that T-bills (liquid) and T-notes (less liquid) of very similar characteristics exhibit average yield spreads of at least 35 basis points. A similar analysis is performed by Boudoukh and Whitelaw (1991) on the benchmark

Japanese government bonds (liquid) and other Japanese bonds (less liquid) which showed the yield spread averages more than 50 basis points.

To understand the purpose of this paper and the model created, a formal definition of liquidity risk must be outlined. Liquidity risk as created in this paper refers to the difference between the quoted market price and the price that a trade of a given size can be executed at. This definition assumes that a transaction of any size can be completed but the asset price will be impacted accordingly based on a liquidity parameter. However, this model also introduces a restriction on the amount of trading that can occur in any period. This definition allows markets to have extreme liquidity costs, while also placing a firm restriction on maximum transaction size as in Longstaff (2001).

The model I will setup is meant to parallel the problem a fund manager might face in light of a capital requirement. If the manager is aware a certain level of capital is required the manager will wish to hedge and trade using optimal strategies, allowing the manager to maximize the expected value of liquidating a position in an asset (this is analogous to minimizing the 'liquidity cost' associated with transactions that impact market prices).

I draw on components from the model used in Cetin, Jarrow, Protter and Warachka (2006) (CJPW). In that paper the asset price follows a geometric Brownian motion and then the approximation created by Cox, Ross and Rubenstein (1979) is applied which reduces the model to a binomial tree structure. In this paper I diverge from the setup by applying different limitations on the continuous process. I start by setting the time step equal to 1 day and the model time to span 5 dates. I then bound holdings in the risky asset by only allowing them to take on integer values within a specified range. I will also derive a static hedge that will be created in the initial period. The optimal discrete trading strategy will then be found using simulation techniques.

Section 2 outlines the model and its parameters and discusses why certain specifications are made. Section 3 provides an explanation of the simulation technique to be used. Section 4 outlines the values used for all variable and static parameters in each simulation and summarizes the results while Section 5 concludes.

2. Model Specification

The economy is composed of three assets. One is a risky asset with holdings at any time, $t \in [0, T]$, defined by X_t. The second asset in the economy is a riskless bond with holdings at any time, $t \in [0, T]$, defined by Y_t. The third asset is a European put where the underlying is the risky asset and the holdings of the put are defined by H_K where K is the exercise price of the option. No time identification is necessary for the put since it is purchased at the initial date and held until maturity in the terminal period. This put is intended to hedge the risky asset. The price of the put is defined by p_K. The risky asset will evolve according to a geometric Brownian motion defined by the following stochastic differential equation:

$$dS = \mu S dt + \sigma S dZ \quad \text{(Neftci (2000))}$$

This has the solution:

$$S_t = S_{t-1} e^{\left(\left(\mu - \left(\frac{1}{2}\right)\sigma^2\right)dt + \sigma \sqrt[2]{t} Z\right)}$$

Where Z is a standard Brownian motion, σ is the daily volatility, μ is the mean daily return and S is the price of the asset (Neftci (2000)). This setup is similar to that used in

CJPW. This process provides paths with a strong resemblance to historical paths of stock prices. To extend the model to analyze assets with price paths exhibiting characteristics different from stocks the stochastic differential equations should be change to another form such as a mean-reverting process (Neftci (2000)). To do so the remaining model specifications may need to be adjusted slightly but it is a fairly simple modification.

This economy is composed of many agents; however, they exist only to be the counterparty on every transaction. The result of this statement is the model is considered a reduced-form illiquidity model as discussed in CJPW because we only consider the liquidity costs created by the transactions initiated by a single agent referred to as simply 'The agent'. The agent acts according to preferences defined by a utility function, $U(\cdot)$. The agent's utility follows the assumptions:

1. $U(\cdot)$ is strictly increasing,

2. $U(\cdot)$ is quasi-concave,

3. $U(\cdot)$ is twice differentiable and

4. $U(\cdot)$ is continuous

The agents utility is a function of the value of terminal wealth, $W_{T\omega}$. $W_{T\omega}$ is the value of the risky asset liquidated in period T plus the bond and its accumulated interest plus the payoff associated with the put. It is now possible to set out the objective function of the agent. The agent seeks to maximize expected utility defined by:

$$E[U(W_{T\omega})] = \int_{0}^{\infty} U(W_{T\omega})$$

This continuous time representation allows us to see what an agent would do to minimize the cost of liquidity. Under continuous time the agent can trade at infinitesimally small intervals. This is very advantageous when facing illiquidity because it gives the agent the capability to trade infinitesimally small amounts an infinitely many times at infinitesimally small intervals. The result of this trading strategy is that the price impact of every transaction is negligible and liquidity costs will converge to zero (CJPW). Initially this seems useful because liquidity costs are annihilated, but this strategy assumes markets are frictionless, which is not the case in reality, since transaction costs exist. Transaction costs will converge to infinity when this strategy is implemented completely wiping out any advantage gained. This leads to a departure from treating the model as continuous by using restrictions. Time is restricted to be discrete by creating a time step of 1 day between each trading period. The model is restricted to a discrete set of states ω where $\omega \in [1, \Omega]$ and Ω is finite. A given state ω is now further specified to represent a given price path created by $S_t = S_{t-1}e^{\left(\left(\mu - \left(\frac{1}{2}\right)\sigma^2\right)dt + \sigma \sqrt[2]{t}Z\right)}$ and each node on the path is denoted by ωt. Expected utility can now be defined as a discrete function:

$$E[U(W_{T\omega})] = \sum_{\omega=1}^{\Omega} \pi_{T\omega} U(W_{T\omega})$$

This expectation is dependent on the probability, π, of state ω at time $t = T$. By using this notation the model is extendible to any $\Omega > 0$ states and any $T \geq 0$ times.

A hedge is created using the European put. This hedge will be initiated in the starting period with an exercise date of the terminal period. The hedge will be static to avoid transaction costs. The cost of the hedge will be based on the value of puts in a perfectly liquid market using the Black-Scholes option pricing formula (Jarrow and Turnbull (1999)). The initial price of this asset will be unaffected by market liquidity because this model is treating liquidity constraints as unanticipated shocks and the Black-Scholes pricing result assumes the existence of frictionless markets.

The agent will create transactions by placing an order for the risky asset of amount x_t. The agent can choose between three types of orders; a buy order ($x_t > 0$), a sell order ($x_t < 0$), or a null order ($x_t = 0$). These orders can only take on integer values within a given range and that allows a maximum of 5 units to be bought or sold. Similar notation is used for orders placed on the riskless bond with y replacing x. The order has an impact on the market price based on the size and type of the order and on the liquidity of the risky asset. The impact on price can be thought of as the cost of illiquidity. The size and type of the order is determined by the discrete trading strategies that will be used for simulation while liquidity is a parameter that will be varied to examine its effect on optimal discrete trading strategies.

Illiquidity is created in the risky asset through a supply curve. The supply curve is nonlinear which creates liquidity costs that are a nonlinear function of transaction size (the marginal liquidity cost of trading is increasing). The supply curve has the functional form:

Supply curve = $e^{\alpha x}$ (Cetin, Jarrow and Protter (2002)(CJP))

Examples of the supply curve for various levels of alpha are shown below:

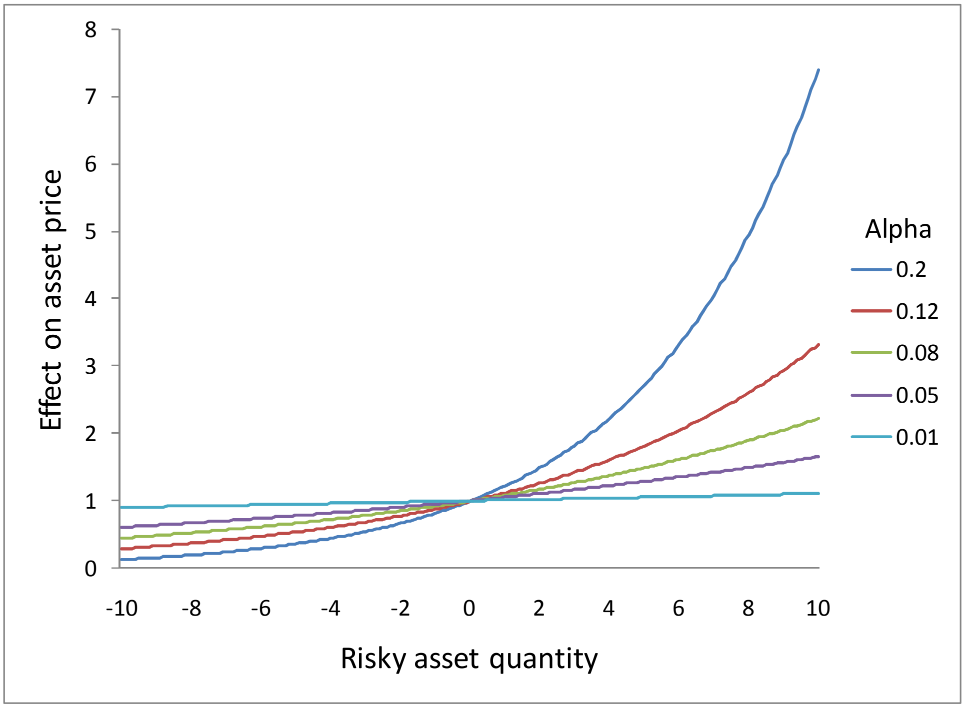

This creates a transaction price similar to that in CJP. The key difference is the inclusion of the persistence of a shock. The resulting transaction price takes the form:

$$S_{t\omega x} = e^{\alpha x}[S_{t\omega 0} - \sum_{i=0}^{t-1} \rho_i^{t-i} c_{liq\omega i}] \, , \forall \, \alpha \geq 0, \forall \, 1 \geq \rho \geq 0$$

This transaction price implies a liquidity cost of:

$$c_{liq\omega t} = \; [S_{t\omega 0} - \sum_{i=0}^{t-1} \rho_i^{t-i} c_{liq\omega i}] - S_{t\omega x}$$

where ρ_i^{t-i} represents the persistence of a liquidity cost from period i on the pre-transaction price of the risky asset at time t. Examples of the persistence parameter in the liquidity cost function are:

i) At $t = 1$ the liquidity cost from $t = 0$ impacts the pre-transaction price at $t = 1$ by $-\rho_0^1 c_{liq\omega 0}$.

ii) At $t = 2$ the liquidity cost from $t = 0$ impacts the pre-transaction price at $t = 2$ by $-\rho_0^2 c_{liq\omega 0}$.

To ensure all liquidity costs are incurred, the agent must fully liquidate all holdings in the final date. This means that if the agent is long (short) a position in the risky asset, a sale (buy) order must be initiated. This leads to a function for terminal wealth defined by:

$$W_{T\omega} = S_{T\omega x} X_{T\omega} + R^T Y_{T\omega} + H_K max\{0, K - S_{T\omega x}\}, \forall \omega \in \mathbb{N}$$

Where $R \equiv (1 + r)$ and r represents the rate on the riskless bond.

The trading strategies used in this model must be self-financing. To satisfy this constraint when a buy (sell) order is initiated for the risky asset a sale (buy) order of equivalent value must be initiated for the riskless bond. The cost of the hedge is financed through a sale of the riskless bonds which will be incorporated into the initial endowments. This ensures that in each period there is a net flow of zero dollars. Based on this, it is possible to create a situation where terminal wealth will take on a negative value. This will be accounted for in the simulation by replacing all negative values with a value of zero to represent the ability to walk away from debt through bankruptcy. This may not be a realistic assumption as the agent will most likely have a lower level of utility for negative terminal wealth (bankruptcy) than the agent would have for a utility derived from zero terminal wealth. This assumption although unrealistic is suitable because it should not provide much distortion in the resulting expected utility levels.

Now that required constraints on trading strategies have been mentioned the optimization problem can be stated mathematically in discrete functional form:

$$MAX\ E[U(W_{T\omega})]$$

$$\Lsh\ MAX \sum_{\omega=1}^{\Omega} \pi_{T\omega} U(W_{T\omega})$$

$$\Lsh\ MAX \sum_{\omega=1}^{\Omega} \pi_{T\omega} U(S_{T\omega x} X_{T\omega} + R^T Y_{T\omega}\ +\ H_K max\{0, K - S_{T\omega x}\})$$

$$s.th.\ \ S_{t1x} x_{t\omega} + R^t y_{t1} + p_0 H_K = 0\ \ ,\ \ t = 0$$

$$and\ \ \ S_{t\omega x} x_{t\omega} + R^t y_{t\omega} = 0\ \ , \forall\ t \in [1, T]\ \ , \forall\ \omega \in \mathbb{N}$$

$$and - 5 \geq x_t \leq 5\ \ , \forall\ t \in [0, T]$$

$$and\ X_t\ can\ only\ take\ on\ integer\ values$$

We are now ready to attempt to solve this problem. The method being used will be simulation.

3. Simulation Approach

A Monte Carlo simulation will be used to find the optimal trading strategy. To make this a computationally feasible problem the trading strategies are being restricted by a set of allowable discrete holding amounts for the risky asset. The discrete trading strategies will then be run on 10000 sample paths and the results analyzed to look at return distribution and the optimal trading strategy. The trading strategy that will result from this simulation represents trading actions that can be undertaken regardless of the stock price that prevails. This means the strategy is only necessarily optimal at $t = 0$ but that it would be the optimal strategy to use based on the information available at $t = 0$. The benefit of this approach is that it provides an investor with two possible actions:

1) The strategy can be created at time 0 and used for all trading decisions. The agent will use the information set available at time 0 to derive a trading strategy. This strategy can be implemented by an investor who is not able to develop a new strategy at every period that would be contingent on the evolution of the price of the risky asset over time. While this strategy is optimal at time 0, it will only be suboptimal at any subsequent period as it only accounts for the realized illiquidity parameter and does not take into account the rest of the new information set available (i.e. the actual price of the risky asset that is realized) and is therefore reliant on a trading decision based on the expected utility from future payoffs and not from any single actual payoff. This strategy is analyzed because it demonstrates what an optimal strategy would be in the initial time when the agent only knows the probability distribution of the illiquidity and prices that will be present in the market in all subsequent periods as opposed to knowing what price actually prevails.
2) At each time, t, the simulation can be rerun using new price (state) information to provide an updated optimal trading strategy that will be contingent on the evolution of the price path. This means the trader would take the new information available at every time period (i.e. price at current and previous times, trades made at previous times and the illiquidity parameter that prevails) and would rerun the simulation with the updated data. This would allow the investor to update the strategy and it would now be contingent on all new information that

was previously unknown. This strategy is created by analyzing the expected utility of the payoffs at the terminal date. However, it should only be thought of as optimal in the period it is derived. This is because the strategy is created from the expected utility, while the actual optimal strategy from a subsequent period onward will be heavily dependant on the price of the risky asset that ends up being realized. The simulations in this paper will not provide results for this strategy as the computations would grow exponentially and the computing time required would be too great.

The simulation will be coded in Matlab. The sample paths will be created using the pseudo-random variable generator built into the program. Using values for the daily mean return and the daily volatility of 7% and 20% respectively, sample paths will be created using the geometric Brownian motion discussed earlier. A few example paths are shown in the following diagram:

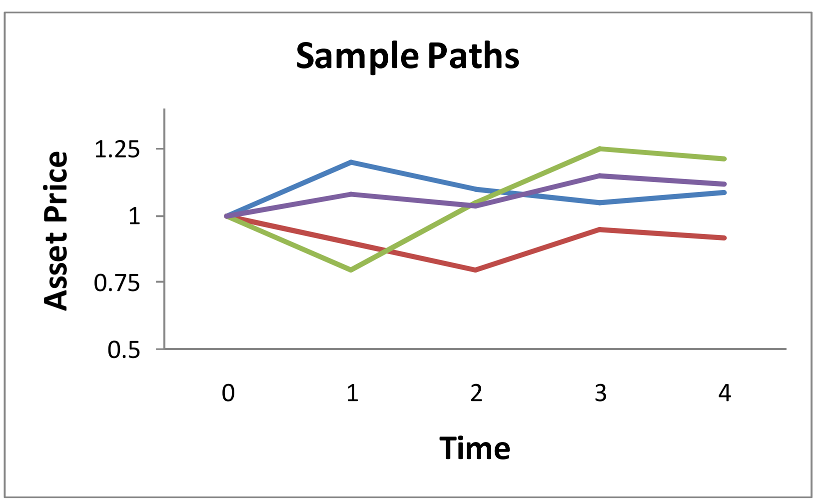

Each sample path will represent a given state ω, have each trading strategy from a pre-specified range input into the model and have the order flows, liquidity costs, and terminal wealth analyzed. The order flows create liquidity costs which modify each asset price path. An example of the modified paths is shown in the following diagram:

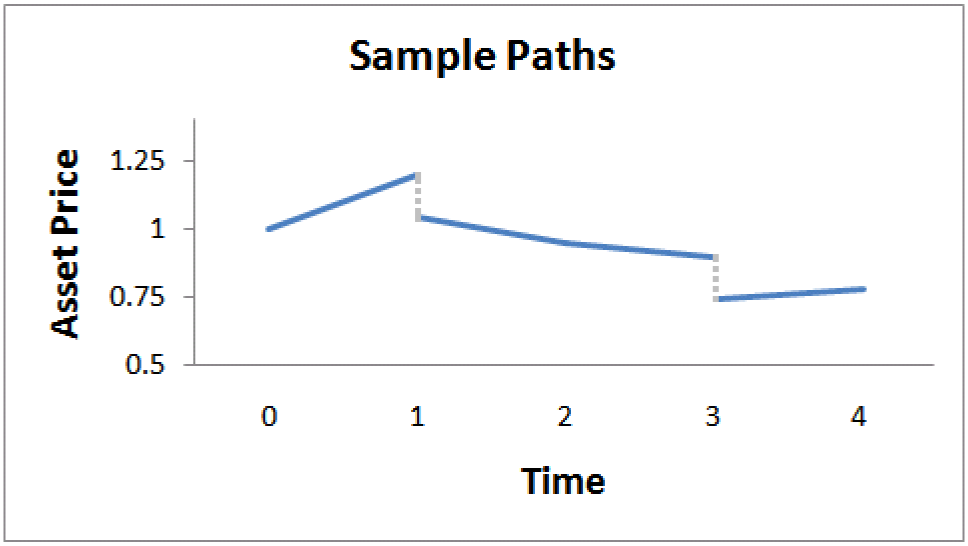

This process yields a given utility value for each state. Each of these values will be collected for the Ω states and the expected utility will be computed. Each path has an equal statistical probability allowing for each path (state) probably to be defined by $\pi_{T\omega} = \frac{1}{\Omega}$ (This definition is possible because of the restriction on Ω to be finite). The utility function will take on one of two forms satisfying the definition of the utility function created earlier. The first form of the function will be linear. This is done to simulate a risk-neutral agent. The second form will be logarithmic. This function is used to simulate a risk-averse agent. These two forms are interesting because of the inclusion of the hedge. The risk-averse agent has preference for safety which should show an increased preference for hedging. All the information required to determine which trading strategy has the highest expected utility has now been collected and it is now possible to determine the optimal strategy for each form of utility.

To examine how optimal trading strategies adjust for assets with differing liquidities, the liquidity parameter, α, will be adjusted. The liquidity parameter will take on values in the range of $\alpha \in [0,0.11]$. The bottom end of the range represents the case where the asset is subject to perfect liquidity and therefore transactions have no effect on price. The top end of the range represents a case where the asset is subject to extreme illiquidity and trades have a large associated liquidity cost. The top end of the range is determined by taking the trade range and letting it be associated with a total price movement of 3-times (Thompson(2003)):

$$e^{\alpha 10} = 3$$

$$\alpha 10 = ln3$$

$$\alpha = \frac{ln3}{10} \cong 0.11$$

This gives a range of $\alpha \in [0,0.11]$. This parameter is meant to act as an unanticipated shock. The shock will occur between time 0 and time 1 with a given probability. The new liquidity parameter will then be present for all remaining periods in the simulation. For example, after the initial period 75% of the paths have perfect liquidity ($\alpha = 0$), 10% have mild illiquidity ($\alpha = 0.01$), 7.5% have moderate illiquidity ($\alpha = 0.03$), 5% have notable liquidity ($\alpha = 0.07$), and 2.5% have extreme illiquidity ($\alpha = 0.11$). All simulations in this paper use the probabilities and illiquidity parameters in the previous example.

The simulation will now be run for differing levels of persistence. The persistence parameter, ρ, will be allowed to take on any value $s.th. 0 \leq \rho \leq 1$. A persistence value of zero represents the case of a purely temporary liquidity cost from a transaction which does not have an impact on prices in subsequent periods. A persistence value of 1 represents the case of a permanent liquidity cost so the cost affects all subsequent periods with a constant dollar impact value.

4. Simulations

In the following sections simulations using varying parameters will be discussed. However, certain parameters will be held constant to allow for comparison between the individual cases. The parameters that will be constant in each simulation are as follows:

1. The mean return on the risky asset (μ) will be equal to 7%.
2. The volatility of the risky asset (σ) will be 20%.
3. The illiquidity factor will take on multiple values within each simulation but will be the same in each case. The illiquidity parameter will take on the following values:

   $$\alpha = [0, 0.01, 0.03, 0.07, 0.11]$$

   The factor will take on the listed values with the following probabilities:

   $$prob = [0.75, 0.1, 0.075, 0.05, 0.025]$$

4. The risk-free rate will be 5%.
5. The initial price of the risky asset and the bond will be 1.
6. The agent will be endowed with 5 units of the risky asset and 5 units of the risk-free asset in period 1.
7. The strike price of the puts on the risky asset will be 1.

8. In scenarios with a hedge the agent will buy 1.2098 puts at a total cost of 0.0857 units and the purchase will be funded through sale of the risk-free asset.

4.1 Simulation 1 – Risk-neutral agent with a static hedge

The first simulation that is run assumes the agent is risk-neutral, liquidity shocks have no persistence ($\rho = 0$) and the portfolio contains a static hedge. This means that the market price of the risky asset is affected by illiquidity only in the period an order is made and not in any other periods and that the agent creates the hedge in the initial period and holds the hedge, as is, until the terminal period. The result of this simulation is summarized in table 1 below where the table values correspond to the holdings of the risky asset.

Risk-Neutral	***alpha (liquidity)***				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	10	10	10	10	10
1	15	10	8	8	8
2	10	8	6	6	6
3	5	4	3	3	3
4	0	0	0	0	0
EU	13.4774	12.743	12.0564	10.9152	9.9047

Table 1: Risk-neutral agent with a static hedge, persistence=0

In the initial period, time 0, the agent buys the maximum amount of stock allowed, which is 5 units, to optimize the utility. The agent is now holding 10 units of the risky asset in the initial period (time=0) for all levels of illiquidity. This decision is made with incomplete information since the agent is unaware which state of liquidity will prevail in period 1. The liquidity shock occurs after the agent has made the decision in the initial period and leads us to five different scenarios, one for each state of liquidity. Depending on which state of liquidity prevails the agent now has five distinct strategies that can be implemented to optimize utility.

In the case of perfect liquidity the agent will buy another 5 units of the risky asset at time 1 (for a total holding of 15 units) and from then on will sell 5 units each period to meet the requirement that in the terminal period the agent's position in the risky asset will be fully liquidated. In this case it is easy to see that the agent is not subject to the liquidity costs, however the agent is still constrained to sell/buy a maximum of five units a period which brings some illiquidity to the market, although not through orders impacting the market price. Under the constraint of a maximum buy/sell quantity the agent attempts to purchase as much of the risky asset allowed because the investor is indifferent towards risk and gains utility from the higher expected payoff of the risky asset compared to the risk-free bond.

When the market stops exhibiting perfect liquidity, the agent's utility is decreasing in illiquidity and the agent closes out the position in the risky asset by selling a portion of the total holdings in each period as opposed to in a single large trade. As illiquidity in the market increases the agent chooses to smooth the selling process. This is done to minimize the total cost of liquidity and is achieved by selling smaller amounts in each period rather than selling in large blocks. This result is analogous to a result from CJPW where the trader sells an infinitesimally small amount in every period and each period is separated by an infinitesimally small interval so the impact of illiquidity on price is almost eliminated. In reality the result of continuously trading infinitesimally small amounts ignores transaction costs which would approach infinity and eliminate the benefit gained by using the strategy. The model in this paper uses time steps equal to one day to prevent continuous trading from occurring.

For the next scenario persistence of the liquidity cost is included. The persistence of the liquidity cost is set at $\rho = 0.6$. This means that a transaction having an impact on price of 1 unit at time 0 will have and impact on price in time 2 of $0.6x1$. All other parameters from the previous scenario remain unchanged. Table 2 summarizes the results below:

Risk-Neutral	***alpha (liquidity)***				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	10	10	10	10	10
1	15	9	8	8	8
2	10	8	6	6	6
3	5	5	3	3	3
4	0	0	0	0	0
EU	13.4774	12.5525	11.5664	9.9997	8.743

Table 2: Risk-neutral agent with a static hedge, persistence=0.6

The scenario with persistence of $\rho = 0.6$ yields very similar results to the situation where $\rho = 0$. The agent will smooth the selling process as illiquidity is increased; the agent's utility is decreasing in α (illiquidity); and when perfect liquidity prevails, the investor holds as much of the risky asset allowed at every period. When comparing the two scenarios it can be seen that for all illiquidity parameter values such that $\alpha > 0$, the agent's utility is less in the case with persistence. The reason that a higher value of persistence causes less utility is because frictions are being increased in the market which makes it harder for the investor to trade the risky asset and the investor gives up utility in order to execute the required trades.

Now we will examine the most extreme case of persistence in this paper where persistence takes on the value $\rho = 1$.The results of this scenario are summarized in Table 3 below:

Risk-Neutral	***alpha (liquidity)***				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	8	8	8	8	8
1	13	10	7	7	7
2	10	9	6	5	5
3	5	5	3	3	3
4	0	0	0	0	0
EU	13.3557	12.5066	11.5744	10.2272	9.1999

Table 3: Risk-neutral agent with a static hedge, persistence=1

This case gives results that are quite similar to the previous two cases where a risk-neutral agent holds a hedged portfolio and persistence takes on a value of $\rho = 0$ or $\rho = 0.6$. However, when comparing the results of this case to previous scenarios, some interesting differences can be seen. More scenarios have been created to give a clearer understanding of the changes that occur when the level of persistence is varied. The complete summary of these simulations can be found in Appendix 1 (a-g). Tables a through g in Appendix 1 show the results of cases that have the same parameters except the level of persistence is varied from $\rho = 0$ to $\rho = 1$.

The first important takeaway from these simulations is that the optimal holding of the risky asset in the initial period decreases in the level of persistence. As persistence is increased the cost of illiquidity in a period has a larger effect on the price level in following periods. This means that by selling today the investor will not only incur costs on the current transaction, but will also incur costs caused by this transaction in all subsequent periods. Since the persistence of illiquidity costs are known, the agent will choose to buy less of the risky asset in the initial period to prevent suffering more illiquidity costs in a later time period. To make this situation preferred, the agent's expected utility must be greater from holding less of the risky asset. For this to happen the risk-free asset must be offering a rate of return that is larger than the expected return on the risky asset less illiquidity costs, otherwise the risk-neutral investor would necessarily prefer the risky asset.

The second important takeaway from these simulations is that for $\alpha > 0$, as ρ increases it eventually becomes possible for the investor to achieve higher utility in the terminal period. This is an important result because it shows that the investor can take advantage of the persistence of a shock to increase utility compared to situations where illiquidity costs are less persistent. As the persistence is increased the investor knows its actions create a longer lasting effect on the market price. The investor can use this knowledge to create a strategy that manipulates the market price in order to achieve higher utility. The following chart shows the expected utility of the agent plotted against the persistence of the illiquidity cost.

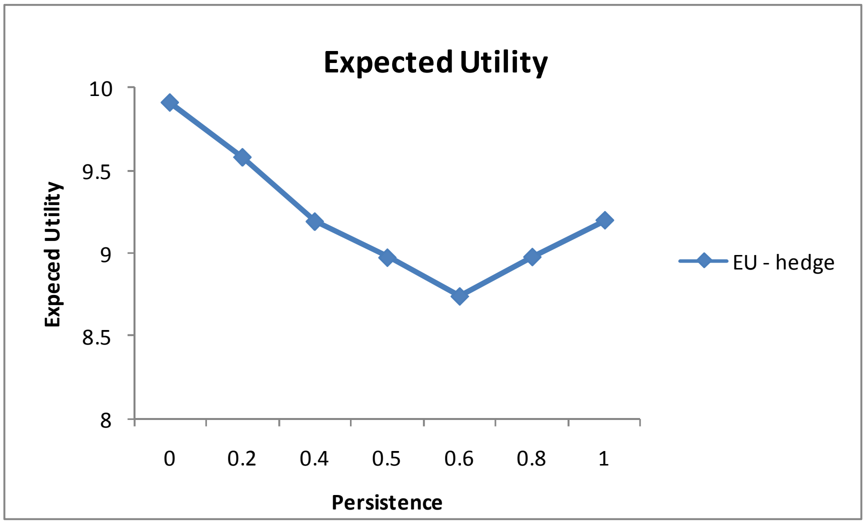

Chart 1: Expected Utility - α=0.11 – Risk-neutral with a static hedge

As seen in the chart above the agent has the least utility when persistence is equal to $\rho = 0.6$. At any point below this level the investor prefers less persistence and at any point above the investor prefers a higher level of persistence.

4.2 Simulation 2 – Risk-neutral agent with no hedge

The first simulation assumes the agent is risk-neutral, liquidity shocks have varying persistence ($\rho = [0, 0.01, 0.03, 0.07, 0.11]$) and the portfolio contains a static hedge. This simulation will assume the agent does not hedge. The scenario with no hedge returned very similar results as the hedged portfolio and the results are summarized in Appendix 2. First, the agent demonstrated a preference to smooth the selling process as illiquidity increased in an attempt to minimize total liquidity costs. Second, the utility of the agent decreases as illiquidity increased. Lastly, in the case of perfect liquidity the investor holds as much of the risky asset that is allowed at every period. The most interesting result of the no hedge case can be identified from the table below:

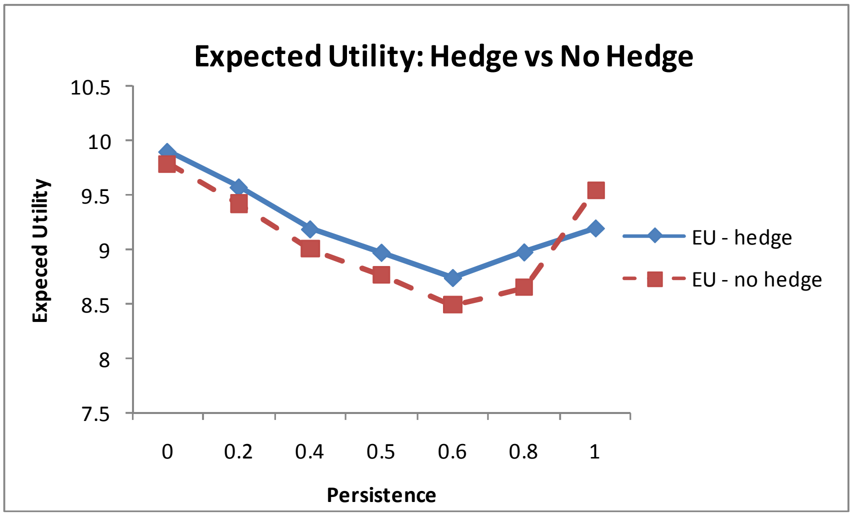

Chart 2: Expected Utility - α=0.11 – Risk-neutral static hedge vs. no hedge

Chart 2 shows that for a level of illiquidity, when $\alpha > 0.01$, the agent benefits from hedging the portfolio. The expected utility of the agent is higher in all cases except when persistence is equal to one. This is because the hedge protects the investor from some of the downside risk which is higher than in a frictionless market because of the illiquidity present. The protection created by the hedge benefits the investor by providing risk reduction and higher expected terminal wealth.

This result is also very interesting because it demonstrates two key properties of the Black-Scholes option pricing formula that was used to price the put in this paper. The first is that Delta hedging (Jarrow and Turnbull (1999)) only works for very small price movements while the second is that the market with illiquidity has frictions and is therefore not complete. These properties can be seen by the fact that the expected utility of the hedge and no hedge scenarios differ, which should not happen since the agent is a risk-neutral investor.

The delta hedging technique will provide an adequate hedge when the underlying asset has a minimal change in price. However, as the price change became larger the initial hedge no longer provides the investor with the appropriate ratio of puts to the risky asset. In order to keep the portfolio properly hedged the investor would need to continually rebalance the position in the option and would incur excessive transaction costs by doing so.

The second property is that the Black-Scholes option pricing formula assumes that markets are complete, transaction size is unbounded and assets can be traded continuously (Jarrow and Turnbull (1999)). This is because the cost of the hedge should fully incorporate the possible price movements in the risky asset and, intuitively, since the agent is risk-neutral only the expected payoff should matter. In this market the price of the put was not able to account for the increased downward pressure on the risky asset and the resulting price of the put was too low, benefiting the investor by yielding higher utility. This demonstrates that if an investor is facing illiquidity it is optimal to have a hedged portfolio.

4.3 Simulation 3 – Risk-averse agent with a static hedge

The first two simulations assumed the agent is risk-neutral. This simulation will assume the agent is risk-averse with a utility function that is logarithmic. This option is explored because it means the investor has preferences that are concave and exhibit decreasing marginal utility of wealth. These preferences are important as it means the investor has an increased relative preference for the risk-free asset. Table 4 below summarizes the results of the simulation where the agent is risk-averse, liquidity shocks have no persistence ($\rho = 0$) and the portfolio contains a static hedge.

Risk-Averse	***alpha (liquidity)***				
Time	**0**	**0.01**	**0.03**	**0.07**	**0.11**
0	5	5	5	5	5
1	7	5	4	4	4
2	9	4	3	3	3
3	4	2	2	2	2
4	0	0	0	0	0
EU	2.5301	2.5152	2.5013	2.4769	2.4539

Table 4: Risk-averse agent with a static hedge, persistence=0

This simulation returned results that have some similarities with the previous scenarios. First, the agent demonstrated a preference to smooth the selling process as illiquidity is increased in order to minimize total liquidity costs. Second, the utility of the agent decreases as illiquidity is increased. However, the important difference between the results of this simulation and the previous one with a risk-neutral agent is that the optimal decision for the investor in the initial period is to make no trade at all and to hold all 5 endowed units of the risky asset until the second period. This is an important outcome as it demonstrates the agent's preference for the risk-free asset.

Examining the case where perfect liquidity prevails after the initial period, it can be seen that the agent will choose to buy 2 more units of the risky asset in the second period followed by another 2 in the third. After this point, the agent is required to start selling the risky asset so that full liquidation can occur in the final period. This means that when the market is perfectly liquid the agent will choose to hold the risky asset to benefit from the higher expected return while not holding the maximum amount allowed. However, it also demonstrates the agent is risk averse because it is optimal for the agent to continue to hold some of the risk-free asset instead of the risky asset to get the guaranteed return.

Simulations for the risk-averse agent with a hedged portfolio will now be examined for varying levels of persistence. The results of the simulations are summarized in Appendix 3. These simulations return very similar results as the case of a risk-neutral agent with a hedged portfolio. As in the risk-neutral simulation, the agent has a preference to hold less of the risky asset in the initial period as persistence is increased. Another similar result can be seen in the chart below:

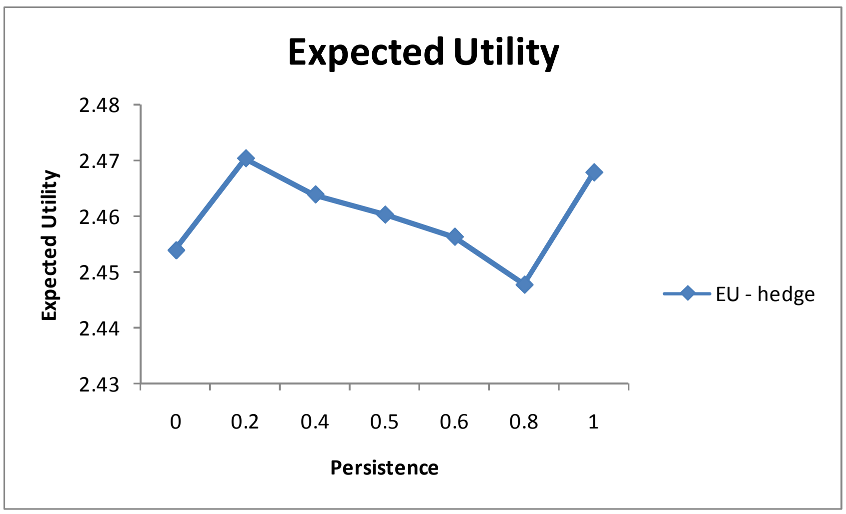

Chart 3: Expected Utility - α=0.11 – Risk-averse with a static hedge

It can be seen in the chart that as the persistence is increased there are two distinct levels that have higher expected utility than the previous level of persistence. This occurs at each of these points the agent chooses a different strategy in the case with higher persistence. This shows that the agent is able to use the knowledge of the persistence of a shock to create a better trading strategy and increase utility.

5. Conclusion

This paper studied a situation in which an investor is trying to maximize the expected utility derived from having a given level of wealth at the end of a series of trades. The investor was forced to trade in an environment that ranged from mildly to extremely illiquid, causing the trades to have an effect on market price.

The simulations in this paper return interesting results that are dependent on the level of illiquidity present in the market and also on the level of persistence of the illiquidity costs. The key results of the simulations are:

1) For any level of persistence, the investor will smooth liquidation as illiquidity is increased.
2) For any level of persistence, as illiquidity increases the expected utility of the agent is decreasing. This is a result of more frictions being present in the market causing the agent not to be able to trade freely.
3) For all levels of illiquidity, as persistence is increased less of the risky asset will be held in the initial period.
4) For all levels of illiquidity greater than $\alpha = 0$, as the level of persistence is increased it eventually becomes possible, as persistence becomes high enough, to increase expected utility compared to lower levels of persistence. This demonstrates that the agent appears to have increased power to manipulate the market as the level of persistence is increased.

The model that was created for this paper could be extended in the following ways. First, the number of dates and the size of the trading range could be increased. This would allow the investor to have more flexibility with the trading strategies. By putting these changes in, the computing power required to run the simulations would increase exponentially, which was a reason for choosing the levels of the parameters in this model. Second, transaction costs could be added that have to be paid every time a trade is executed. By adding the transaction costs it would create a deterrent to trading continuously because every time a trade is made the agent would incur a cost. Finally, competition could be introduced into the model. By including the activities of another agent it would create a situation where the price of the risky asset would be affected by both agents. This makes the optimal strategy of both agents dependent on each other. Under this last proposed extension the optimal trading strategy would become highly correlated with the difference between the agents' preferences and endowments.

A. Appendices

A.1 Risk-neutral agent with a static hedge

a)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.00		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	*alpha (liquidity)*				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	10	10	10	10	10
1	15	10	8	8	8
2	10	8	6	6	6
3	5	4	3	3	3
4	0	0	0	0	0
EU	13.4774	12.743	12.0564	10.9152	9.9047

b)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.20		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	*alpha (liquidity)*				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	10	10	10	10	10
1	15	9	8	8	8
2	10	7	6	6	6
3	5	4	3	3	3
4	0	0	0	0	0
EU	13.4774	12.6799	11.9254	10.6631	9.5767

c)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.40		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	*alpha (liquidity)*				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	10	10	10	10	10
1	15	9	8	8	8
2	10	7	6	6	6
3	5	4	3	3	3
4	0	0	0	0	0
EU	13.4774	12.6197	11.7638	10.3602	9.1912

d)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.60		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	*alpha (liquidity)*				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	10	10	10	10	10
1	15	9	8	8	8
2	10	8	6	6	6
3	5	5	3	3	3
4	0	0	0	0	0
EU	13.4774	12.5525	11.5664	9.9997	8.743

e)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.80		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	alpha (liquidity)				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	9	9	9	9	9
1	14	9	7	7	7
2	10	8	5	5	5
3	5	5	3	3	3
4	0	0	0	0	0
EU	13.4165	12.5132	11.551	10.1123	8.9796

f)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	1.00		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	alpha (liquidity)				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	8	8	8	8	8
1	13	10	7	7	7
2	10	9	6	5	5
3	5	5	3	3	3
4	0	0	0	0	0
EU	13.3557	12.5066	11.5744	10.2272	9.1999

A.2 Risk-neutral agent with no hedge

a)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.00		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	*alpha (liquidity)*				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	10	10	10	10	10
1	15	10	8	8	8
2	10	8	6	6	6
3	5	4	3	3	3
4	0	0	0	0	0
EU	13.5011	12.7453	12.0499	10.859	9.7895

b)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.20		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	*alpha (liquidity)*				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	10	10	10	10	10
1	15	9	8	8	8
2	10	7	6	6	6
3	5	4	3	3	3
4	0	0	0	0	0
EU	13.5011	12.6885	11.911	10.5856	9.4265

c)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.40		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	*alpha (liquidity)*				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	10	10	10	10	10
1	15	9	8	8	7
2	10	7	6	6	5
3	5	4	3	3	3
4	0	0	0	0	0
EU	13.5011	12.6249	11.7379	10.252	9.0064

d)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.60		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	*alpha (liquidity)*				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	10	10	10	10	10
1	15	9	8	7	7
2	10	8	6	5	5
3	5	5	3	3	3
4	0	0	0	0	0
EU	13.5011	12.5512	11.5236	9.8504	8.491

e)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.80		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	*alpha (liquidity)*				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	9	9	9	9	9
1	14	9	7	7	7
2	10	8	5	5	5
3	5	5	3	3	3
4	0	0	0	0	0
EU	13.4402	12.5094	11.4942	9.9221	8.6498

f)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	1.00		- risky	1.00		- risky	5
Risk-neutral	- Linear Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Neutral	***alpha (liquidity)***				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	7	7	7	7	7
1	12	9	6	6	6
2	10	8	5	4	4
3	5	5	3	2	2
4	0	0	0	0	0
EU	13.3185	12.534	11.7298	10.517	9.5414

A.3 Risk-averse agent with a static hedge

a)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.00		- risky	1.00		- risky	5
Risk-averse	- Logarithmic Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Averse	***alpha (liquidity)***				
Time	**0**	**0.01**	**0.03**	**0.07**	**0.11**
0	5	5	5	5	5
1	7	5	4	4	4
2	9	4	3	3	3
3	4	2	2	2	2
4	0	0	0	0	0
EU	2.5301	2.5152	2.5013	2.4769	2.4539

b)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.20		- risky	1.00		- risky	5
Risk-averse	- Logarithmic Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Averse	***alpha (liquidity)***				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	4	4	4	4	4
1	7	4	3	3	3
2	9	4	2	2	2
3	4	2	1	1	1
4	0	0	0	0	0
EU	2.5299	2.5144	2.5027	2.486	2.4702

c)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.40		- risky	1.00		- risky	5
Risk-averse	- Logarithmic Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Averse	***alpha (liquidity)***				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	4	4	4	4	4
1	7	4	3	3	3
2	9	4	2	2	2
3	4	2	1	1	1
4	0	0	0	0	0
EU	2.5299	2.5138	2.5006	2.4814	2.4638

d)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.60		- risky	1.00		- risky	5
Risk-averse	- Logarithmic Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Averse	***alpha (liquidity)***				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	4	4	4	4	4
1	7	4	3	3	3
2	9	4	3	2	2
3	4	2	2	1	1
4	0	0	0	0	0
EU	2.5299	2.5132	2.4982	2.4763	2.4563

e)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	0.80		- risky	1.00		- risky	5
Risk-averse	- Logarithmic Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Averse	***alpha (liquidity)***				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	4	4	4	4	4
1	7	4	4	3	3
2	9	4	3	2	2
3	4	2	2	1	1
4	0	0	0	0	0
EU	2.5299	2.5125	2.4956	2.4701	2.4478

f)

Parameters							
5 date model			Price at time 0			Endowments	
Persistence	1.00		- risky	1.00		- risky	5
Risk-averse	- Logarithmic Utility		- risk-free	1.00		- risk-free	4.9143
Mean-return	7%		- strike price of put	1.00		- put	1.2098
Daily volatility	20%						
Interest rate	5%						

Risk-Averse	***alpha (liquidity)***				
Time	***0***	***0.01***	***0.03***	***0.07***	***0.11***
0	3	3	3	3	3
1	7	4	3	3	3
2	9	4	3	2	2
3	4	2	2	1	1
4	0	0	0	0	0
EU	2.5292	2.5133	2.5004	2.4829	2.4677

References

Amihud, Y., and H. Mendelson, 1991, "Liquidity, Maturity, and the Yields on U.S. Treasury Securities," Journal of Finance, 46, 1411-1425.

Black, F., and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," Journal of Political Economy, 81, 637-654.

Boudoukh, J., and R.F. Whitelaw, 1991, "The Benchmark Effect in the Japanese Government Bond Market," Journal of Fixed Income, September, 52-59.

Brunnermeier, Markus K, 2008, "Deciphering the Liquidity and Credit Crunch 2007-08."

Cetin, U., R. Jarrow, P. Protter, 2002, "Liquidity Risk and Arbitrage Pricing Theory", Finance and Stochastics, 8, 311-341.

Cetin, U., R. Jarrow, P. Protter, M. Warachka, 2006, "Pricing Options in an Extended Black-Scholes Economy with Illiquidity: Theory and Empirical Evidence," Review of financial studies, 19, 493-529.

Cox, J., and C.-F. Huang, 1989, "Optimal Consumption and Investment Policies when Asset Prices Follow a Diffusion Process," Journal of Economic Theory, 49, 33-83.

Cox, J., S. Ross, M. Rubinstein, 1979, “Option Pricing: A Simplified Approach,” Journal of Financial Economics, 7, 229-263.

Jarrow, R.A., and S.M. Turnbull, 1999, “Derivative Securities,” Mason: Cengage South-Western

Longstaff, F.A., 2001, “Optimal Portfolio Choice and the Valuation of Illiquid Securities,” The Review of Financial Studies, 14, 407-431.

Lowenstein, R., 2001, “When Genius Failed : The Rise and Fall of Long-Term Capital Management.” New York: Random House Trade Paperbacks, 2001.

Merton, R., 1969, “Lifetime Portfolio Selection under Uncertainty: The Continuous Time Model,” Review of Economics and Statistics, 51, 247-257.

Merton, R., 1971, “Optimum Consumption and Portfolio Rules in a Continuous Time Model,” Journal of Economic Theory, 3, 373-413.

Merton, R., 1973, “The Theory of Rational Option Pricing,” Bell Journal of Economics and Management Science, 4, 141-183.

Neftci, S.N., 2000, “An Introduction to the Mathematics of Financial Derivatives,” New York: Academic Press.

Thompson, J., 2003, “Simulating Liquidity Risk: A Discrete Time Analysis.”