
An Order-Driven Agent-Based Artificial Stock Market to Analyze Liquidity Costs of Market Orders in the Taiwan Stock Market

Yi-Ping Huang¹, Shu-Heng Chen², Min-Chin Hung¹, and Tina Yu³

¹ Department of Financial Engineering and Actuarial Mathematics,
Soochow University, Taipei, Taiwan
berthuang@cht.com.tw, nhungg@scu.edu.tw

² AI-ECON Research Center, Department of Economics,
National Chengchi University, Taipei, Taiwan
chen.shuheng@gmail.com

³ Department of Computer Science, Memorial University of Newfoundland,
St John's, NL A1B 3X5, Canada
gwoingyu@gmail.com

Summary. We developed an order-driven agent-based artificial stock market to analyze the liquidity costs of market orders in the Taiwan Stock Market (TWSE). The agent-based stock market was based on the DFGIS model proposed by Daniels, Farmer, Gillemot, Iori and Smith [2]. When tested on 10 stocks and securities in the market, the model-simulated liquidity costs were higher than those of the TWSE data. We identified some possible factors that have contributed to this result: 1) the overestimated effective market order size; 2) the random market orders arrival time designed in the DFGIS model; and 3) the zero-intelligence of the artificial agents in our model. We continued improving the model so that it could be used to study liquidity costs and to devise liquidation strategies for stocks and securities traded in the Taiwan Stock Market.

9.1 Introduction

Market liquidity, or the ability of an asset to be sold without causing a significant amount of price movement and with minimum loss of value, plays an important role in financial investment and in securities trading. One recent event that highlighted the impact of asset liquidity on financial institutions was the collapse of Bear Stearns during the subprime crises. Bear Stearns was involved in securitization and issued a large amount of asset-backed securities, mostly mortgage-backed assets, which fell dramatically in liquidity and value during the crisis. In March 2008, the Federal Reserve Bank of New York provided an emergency loan to try to avert the collapse of the company. However, the company could not be saved and was subsequently sold to JP Morgan Chase in 2008.

In large investment institutions, the liquidation of a large block of assets within a given time constraint to obtain cash flow arises frequently. For example, a financial institution may

need to liquidate part of its portfolio to pay for its immediate cash obligations. One possible liquidation strategy is to sell the entire block of assets at once. However, this high-volume trading can cause the price of the share to drop between the time the trade is decided and the time the trade is completed. This implicit cost (due to the price decline) is known as the market impact cost (MIC) or liquidity cost (the numerical definition is given in Section 9.3). To minimize such cost, a better strategy is to divide the block of assets into chunks and sell them one chunk at a time. However, in what way should those chunks be sold so that the liquidity cost is minimized?

In Algorithmic Trading, where computer programs are used to perform asset trading including deciding the timing, price, or the volume of a trading order, this liquidation problem is characterized as an optimization problem. With a smooth and differentiable utility function, the problem can be solved mathematically [1, 6].

However, this mathematical approach to find an optimal liquidation strategy has some shortcomings, such as the imposed assumption that price dynamics are an arithmetic random walk with independent increments. In this chapter, we adopt a different approach by devising an agent-based artificial stock market, which is empirically driven and simulation-based (the model is explained in Section 9.3). It, therefore, can in principle work with various stochastic processes where analytical solutions would be difficult to derive. By performing simulations and analyzing liquidity costs induced under different market scenarios, we hope to understand the dynamics of liquidity costs, and hence to devise a more realistic optimal liquidation strategy.

The rest of this chapter is organized as follows. In Section 9.2, we provide the background and summarize related works. Section 9.3 explains the agent-based artificial stock market we developed based on the DFGIS model and based on data from the Taiwan Stock Market (TWSE). In Section 9.4, the 10 securities and stocks that we selected to conduct our study are presented. Section 9.5 provides the model parameters used to perform our simulation. We analyze the simulation results in Section 9.6 and present our discussions in Section 9.7. Finally, Section 9.8 concludes the chapter with an outline of our future work.

9.2 Background and Related Work

This study implements an agent-based model for an order-driven double auction market, which is the most common financial market in the world. We shall first give a brief introduction to the basic microstructure and trading mechanism of a standard order-driven double-auction market. Next, the DFGIS model [2] on which our agent-based artificial stock market is based will be presented. After that, we will summarize the work of [5] on agent-based models used to study liquidation strategies at the end of the section.

9.2.1 Order-Driven Double-Auction Markets

In an order-driven double-auction market, prices are determined by the publication of *orders* to buy or sell shares. This is different from a quote-driven market where prices are determined from *quotations* made by market makers or dealers.

There are two basic kinds of orders in an order-driven market. Impatient traders submit *market orders*, which are requests to buy or sell a given number of shares immediately at the best available price. More patient traders submit *limit orders*, which specify the limit (best acceptable) price for a transaction. Since limit orders often fail to result in immediate transactions, they are stored in a *limit order book*. As shown on the left of Table 9.1, limit

Table 9.1 An example of an order book before (left) & after (right) a transaction

Limit Buy Orders		Limit Sell Orders		Limit Buy Orders		Limit Sell Orders	
Size	Price	Size	Price	Size	Price	Size	Price
20	\$1.10	35	\$1.17	15	\$1.09	35	\$1.17
25	\$1.09	10	\$1.19	35	\$1.05	10	\$1.19
35	\$1.05	20	\$1.21	10	\$1.01	20	\$1.21
10	\$1.01	15	\$1.25			15	\$1.25

buy orders are stored in decreasing order of limit prices while limit sell orders are stored in increasing order of limit prices. The buy limit orders are called *bids* and the sell limit orders are called *asks*. For a normal double-auction market, the best (highest) bid price is lower than the best (lowest) ask price. The difference between the two is called the *spread* of the market. In the example in Table 9.1 (left), the *spread* is \$0.07.

When a *market order* arrives, it is matched against limit orders on the opposite side of the book. For example, when a market sell order for 30 shares arrives at the market whose order book is as that on the left of Table 9.1, it will first be matched against the current best bid (20 shares at \$1.10 per share). Since the size of the sell order (30 shares) is larger than that of the best bid (20 shares), the remainder of the market sell order (10 shares) will be matched against the next best bid (25 shares at \$1.09 per share). After the transaction is completed, the limit order book will change to the right of Table 9.1 and the market spread will widen to \$0.08.

9.2.2 The DFGIS Model

In the original DFGIS model [2], all the order flows (including limit orders and market orders) are modeled as a Poisson process. Market orders arrive at the market in chunks of σ shares (where σ is a fixed integer) at an average rate of μ per unit of time. A market order may either be a buy market order or a sell market order with equal probability.

Limit orders arrive at the market in chunks of σ shares, at an average rate of α shares per unit price per unit of time. A limit order may either be a limit buy order or a limit sell order with equal probability. The limit prices in limit orders are generated randomly from a uniform distribution. In particular, the limit sell prices have a range between $(-\infty, a(t))$, where $a(t)$ is the best (lowest) *ask* price in the market at time t . Similarly, the limit buy prices have a range between $(b(t), \infty)$, where $b(t)$ is the best (highest) *bid* in the market at time t . The prices are first converted to logarithms so that the value is always positive. In addition, the price changes are not continuous, but have discrete quanta called *ticks* (represented as dp). Tick size is the price increment/decrement amount allowed in a limit order.

DFGIS also allows the limit order to expire or to be cancelled after being placed in the market. Limit orders are expired and cancelled according to a Poisson process, analogous to radioactive decay, with a fixed-rate δ per unit of time. Table 9.2 lists the parameters of the DFGIS model.

To keep the model simple, the DFGIS does not explicitly allow limit orders whose prices cross the best bid price or the best ask price. In other words, the price of a limit buy order must be below the best ask price and the price of a limit sell order must be above the best bid price. Farmer, Patelli and Zovko [4] implemented the model to explicitly handle this type of order. They defined *effective market orders* as shares that result in transactions immediately and *effective limit orders* as shares that remain on the order book. A limit order with a price

Table 9.2 Summary of DFGIS parameters

Parameter	Description	Dimension
α	avg. limit order rate	share/price · time
μ	avg. market order rate	share/time
δ	avg. limit order decay rate	1/time
σ	order size	a constant share
dp	tick size	price

Table 9.3 An example of an order book before (left) & after (right) handling a limit sell order with price equal to the best bid

Limit Buy Orders		Limit Sell Orders		Limit Buy Orders		Limit Sell Orders	
Size	Price	Size	Price	Size	Price	Size	Price
20	\$1.10	35	\$1.17	25	\$1.09	10	\$1.10
25	\$1.09	10	\$1.19	35	\$1.05	35	\$1.17
35	\$1.05	20	\$1.21	10	\$1.01	10	\$1.19
10	\$1.01					20	\$1.21

that crosses the opposite best price is split into effective market orders and effective limit orders according to the above definition. For example, when a limit sell order of 30 shares with price \$1.10 arrives at the market (with the order book as that listed in Table 9.3 (left)), the order will be split into an effective market order of 20 shares and an effective limit sell order of 10 shares. After the execution of the 20 shares of the effective market order (at price \$1.10), the order book is changed to that on the right of Table 9.3.

9.2.3 The Guo Agent-Based Stock Market Model

Guo [5] implemented an agent-based artificial stock market based on the DFGIS model (he called it the SFGK model) to study time-constrained asset liquidation strategies through market sell orders only. In particular, he compared the performance of two strategies. The first one uniformly divides the liquidation shares X and the time constraint T into N chunks. This “uniform rhythm” strategy instructs a trader to sell X/N shares every T/N seconds, regardless of the market condition.

The second strategy is the “non-uniform rhythm” strategy which also divides the liquidation shares and time uniformly. However, within each time segment, this strategy requires a trader to continuously observe the market spread and initiates the selling of the X/N shares as soon as the current market spread, for the first time within the time segment, falls below the pre-determined spread threshold. If the market spread never falls below the spread threshold for a time segment, the strategy will involve selling the X/N shares at the end of the time segment.

Guo devised one agent (A) with the “uniform rhythm” strategy and another agent (B) with the “non-uniform rhythm” strategy. He tested each agent individually by running 200 simulations independently. For each simulation run, the number of liquidation shares (X) is 20 and the time constraint (T) is 5 minutes. The total asset is divided into 10 chunks, each of which contains 2 shares that will be sold within a time segment of 30 seconds. For each market order the agent submits, the number of shares (α) is 1.

The performance of the two agents is evaluated by the *average selling price per share* relative to the *volume weighted average price in the market*. In other words, it is the measure of how much better or worse an agent performs, when tested compared to the market. His simulation results indicated that agent *B* (based on the “non-uniform rhythm” strategy) outperformed the market while agent *A* (based on the “uniform rhythm” strategy) underperformed the market.

Guo did not explicitly study liquidity costs when devising his liquidation strategy. By contrast, we are interested in quantifying the cost in a real-world stock market. We therefore used TWSE stock order and transaction data (see Section 9.4) to construct the agent-based artificial stock market using the DFGIS model. We describe this agent-based system in the following section.

9.3 An Agent-Based Taiwan Stock Market Model

The agent-based artificial stock market consists of zero-intelligent agents who place *buy*, *sell* or *cancellation* orders at random, subject to the constraints imposed by the current prices. The distribution of the order prices and quantities and the distribution of the time intervals of the order submissions in the market follow that of the DFGIS model. Unlike the Algorithmic Trading model, which imposes unrealistic assumptions, the agent-based model is governed by the 5 DFGIS model parameters. The market properties emerge from the stochastic simulation. We describe our implementation of the abstract DFGIS model for the TWSE in the following subsections.

9.3.1 Buy and Sell Orders

On the TWSE, a submitted order (either buy or sell) needs to specify the price, in addition to the quantity, that the trader is willing to accept. However, the price can cross the *disclosed best prices* (explained in Section 9.3.3). When the price of a buy order is greater than or equal to the disclosed best ask price or the price of a sell order is less than or equal to the disclosed best bid price, there is a match for the transaction. We follow [4] (see Section 9.2.2) and call the portion of an order that results in an immediate transaction an *effective market order*. The non-transacted portion of a submitted order is recorded on the order book and is called the *effective limit order*.

9.3.2 Event-Time Model

We implemented the artificial stock market as an event-time model, where the events in the model are not connected to real time. We first partitioned a trading day into a fixed number of time intervals. The events taking place at each time interval become the event-time series describing the market activities of that day.

The TWSE opens at 9:00 am and closes at 1:30 pm. With a time interval of 0.01 seconds, the event-time series of a trading day is 1,620,000 time intervals long. The TWSE is a *call auction* market where the submitted orders are matched once every 25 seconds. Hence, there are 648 order-matching events in a daily even-time series and the number of events between two order-matching events is 2,499. There are five possible events in an event-time series:

- Effective limit order submission: This has an average rate α and is denoted by L .
- Effective market order submission: This has an average rate μ and is denoted by M .

- Order cancellation submission: This has an average rate δ and is denoted by C.
- Order matching: The buy and sell orders are matched for transactions and is denoted by T.
- No activity: None of the above events occurred in the market and is denoted by N.

From a simulation point of view, the simulation result based on an event-time series and that based on a real-time series are equivalent. For example, the event-time series LNNNM is equivalent to the real-time series L... (real time elapse)... M. However, the event-time model is easier to implement and has a shorter simulation running time because there is no need to handle the time elapse between two events. We therefore adopted the event-time model to implement our artificial stock market.

9.3.3 Order Pricing Rules

As mentioned previously, the TWSE matches submitted orders once every 25 seconds. After a match is completed, some of the orders become effective market orders while some become effective limit orders. The five best prices are then disclosed to the public. During the following 25 seconds of the waiting-for-matching period, the TWSE does not disclose any information about the newly-submitted orders. Consequently, investors do not have the updated best prices, but rather the best prices from the previous matching period to make trading decisions. These *disclosed best prices* are then used to decide order pricing ranges and to calculate the liquidity costs (explained in the next subsection).

The TWSE has a pricing rule whereby the price range of an order has to be between the closing price on the previous trading day (cp) $\pm 7\%$. Thus, the submitted orders in our simulation system have the following price ranges:

- Effective limit buy orders: ($cp-7\%$, $da(t)$), where $da(t)$ is the disclosed best (lowest) *ask* price in the market at time t ,
- Effective limit sell orders: ($db(t)$, $cp+7\%$), where $db(t)$ is the disclosed best (highest) *bid* in the market at time t .
- Effective market buy orders: $cp+7\%$. This guarantees an immediate transaction.
- Effective market sell orders: $cp-7\%$. This guarantees an immediate transaction.

The tick sizes (the price increment/decrement amount) are as defined by the TWSE.¹

9.3.4 Liquidity Costs

We define *liquidity cost* as the difference between the *expected transaction payment* and the *actual transaction payment* of an effective market order. The expected transaction payment is calculated by multiplying the disclosed best price by the number of shares of an effective market order. Since the disclosed best price is from the previous transaction period, this is the expected amount of payment to be made during the current matching period. The actual transaction payment, however, can be different from what was expected. It is calculated as the *executed transaction price* (explained in the next paragraph) multiplied by the the number of shares in a transaction. We can interpret liquidity cost as the difference between an investor's expected transaction payment and the actual transaction payment he/she made.

The TWSE uses a special rule to decide the execution transaction price to match the submitted orders. Instead of the best ask and the best bid in the current matching period, the

¹ See <http://www.twse.com.tw/ch/trading/introduce/introduce.php#1>

Table 9.4 An example of the determination of the execution transaction price

accumulated buy shares	buy shares	price	sell shares	accumulated sell shares	transaction volume (shares)
99	99	102.5	1	103	99
99		102	2	102	99
121	22	101.5	99	100	100
126	5	100	1	1	1

rule selects the price that gives the maximum transaction volume as the execution transaction price. Table 9.4 gives an example of how the transaction price is decided.

In column 2, the buy order quantities under the prices in column 3 are given. Column 4 gives the sell order quantities under the prices in column 3. As an example, the first row shows that there is a buy order bidding \$102.5 for 99 shares and there is a sell order asking for the same price for 1 share. Column 6 gives the number of transaction shares under the price in column 3. In this case, \$101.5 gives the largest transaction volume (100 shares) and is selected as the execution transition price for this matching period.

With the selected execution transaction price (TP) and the disclosed best ask (DBA), the liquidity cost of an *effective market buy order* with volume V is defined as:

$$LC_{buy} = \frac{TP \times V - DBA \times V}{DBA \times V} \tag{9.1}$$

Similarly, with the disclosed best bid (DBB), the liquidity cost of an *effective market sell order* is defined as:

$$LC_{sell} = \frac{DBB \times V - TP \times V}{DBB \times V} \tag{9.2}$$

We normalized the liquidity cost such that the value is the ratio of the original cost to the expected transaction payment.

Both LC_{buy} and LC_{sell} can be positive or negative. This is because, under the TWSE pricing regulation, the execution transaction price may be higher or lower than what a trader has expected. Liquidity cost with a negative value means that the execution transition price is better than what the trader has expected while liquidity cost with a positive value means that the opposite situation applies.

When an order is partially or not executed in the current matching period, the non-executed portion remains in the order book. These effective limit orders may later be executed and become effective market orders. However, effective limit orders may lead to a loss of opportunities related to the changing market prices or a decaying value of the information responsible for the original trading decision. This so-called *opportunity cost* is difficult to estimate, and hence is not considered in our liquidity costs calculation.

9.3.5 Measurement of Model Parameters

The model parameters are estimated using real data from the TWSE (see Section 9.4). These parameters include an effective market order rate (μ), effective limit order rate (α),

Table 9.5 Model Parameter Estimation

Parameter	Description	Value	Dimension
$p(\mu)$	avg. daily effective market order event rate	$\sum_{t=1}^n p_t(\mu) \times w_t$	order-event/time
$p(\alpha)$	avg. daily effective limit order event rate	$\sum_{t=1}^n p_t(\alpha) \times w_t$	order-event/time
$p(\delta)$	avg. daily cancellation order event rate	$\sum_{t=1}^n p_t(\delta) \times w_t$	order-event/time
σ	avg. order size	$\sum_{t=1}^n \sigma_t \times w_t$	share/order

cancellation order rate (δ), and order size (σ). Note that tick size (dp) has been discussed previously in Section 9.3.3.

For each parameter, we calculated the mean of the daily value weighted by the number of daily events. For example, the parameter $p_t(\mu)$ is the ratio of the number of effective market order events (including buy and sell orders) to the total number of buy, sell, cancellation orders and non-active events (1,619,352) on day t . The weight factor w_t is the ratio of the number of order events (including effective market, effective limit and cancellation orders) on day t to the total number of order events for the entire period:

$$w_t = \frac{n_{\mu t} + n_{\alpha t} + n_{\delta t}}{\sum_{i=1}^n n_{\mu i} + n_{\alpha i} + n_{\delta i}} \quad (9.3)$$

where $n_{\mu t}$ is the number of effective market orders on day t ; $n_{\alpha t}$ is the number of effective limit orders on day t ; $n_{\delta t}$ is the number of cancellation orders on day t ; and n is the number of days in the entire period. We measured $w_t \times p_t(\mu)$ across the whole period (77 trading days in this case) and then added them together, which becomes the average daily effective market order event rate $p(\mu)$. Note that its dimension is order-event/time, which is slightly different from μ of [2] whose dimension is share/time (see Table 9.2). We applied the same method to calculate $p(\alpha)$ and $p(\delta)$. With that, we can calculate the average daily non-activity event rate ($p(n)$) as $1 - p(\mu) - p(\alpha) - p(\delta)$.

To calculate the average order size, we first computed the average number of shares in the effective market and effective limit orders submitted on day t (σ_t) (excluding those submitted before the first best prices were disclosed and after the market was closed). The summation of $\sigma_t \times w_t$ for the entire period becomes the average order size σ .

In the simulation, we used a stochastic order size, which is generated randomly from a half-normal distribution with standard deviation $\sqrt{\frac{\pi-2}{2}} \times \sigma$ (σ is the average order size, not standard deviation)[7]. According to [2], this stochastic order size gives the same result as that produced from the constant order size σ . According to the TWSE regulation, the maximum order size is 499,000 shares. Table 9.5 summarizes the model parameters implemented in our system, where n stands for the number of trading days (77) in the data set. Note that this approach to the estimation of the model parameters is similar to that of [4]. It assumes that these parameters are time-invariant, which is an important assumption in this study.

9.3.6 Program Implementation and System Flow

The simulation program was implemented in the Python programming language. Figure 9.1 depicts the overall system workflow. Each simulation is for one trading day for the TWSE.

Initially a series of events on a trading day is generated, based on $p(\mu)$, $p(\sigma)$, $p(\delta)$ and $p(n)$. The number of events is 1,620,000. These events are then executed sequentially, according to what types of events they are.

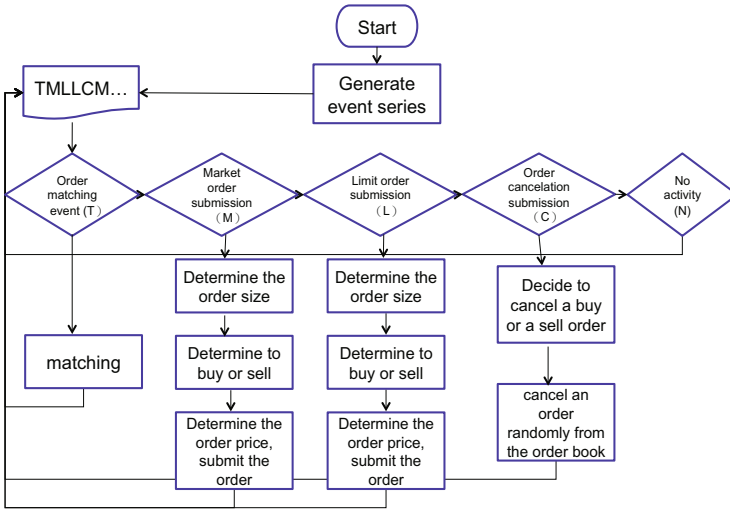


Fig. 9.1 The overall system workflow

If the event is an order matching event (T), the program matches orders and carries out transactions. If it is an effective market order submission (M) or an effective limit order submission event (L), the program decides the order size based on the half-normal distribution of α . Next, the program decides if it is a buy or a sell order with an equal probability (50%). After that, the order price is determined by a *uniformly* random draw from the range given in Section 9.3.3 as what the zero-intelligence agent is supposed to do.² The order is then submitted.

If it is an order cancellation submission event (C), the program decides whether to cancel a buy or a sell order with equal probability (50%). Next, an order on the order book is cancelled randomly. If it is a non-activity event (N), the program continues to process the next event.

9.4 The Data Set

In recent years, the types of securities traded in a stock market have increased from stocks and warrants to Exchange Traded Funds (ETFs) and Real Estate Investment Trust Funds (REITs). ETFs are baskets of stocks which are vehicles for passive investors who are interested in long-term appreciation and limited maintenance. A REIT is a popular investment option as it has better liquidity, in theory, than real estate. We therefore selected a variety of stocks, ETFs and REITs traded in the TWSE to study liquidity costs. They are selected to cover a wide variety of characteristics (see Table 9.6).

² Here, we simply follow [4] to operate the zero-intelligence agent.

Table 9.6 The 10 selected securities and their characteristics

Security	Ticker	Characteristics
Taiwan Top50 Tracker Fund	0050	ETF with the highest trading volume
Polaris/P-shares Taiwan Dividend+ ETF	0056	ETF with a low trading volume
Cathay No.2 Real Estate Investment Trust	01007T	REIT with a high trading volume
Gallop No.1 Real Estate Investment Trust Fund	01008T	REIT with a low trading volume
China Steel	2002	Blue chip stock in TWSE
TSMC	2330	Stock with a high trading volume and the largest market capitalization
MediaTek	2454	Stock with a high unit price
HTC	2498	Stock with a high unit price
President Chain Store	2912	Stock with a large market capitalization but a low trading volume
Inotera	3474	Non-blue chip stock on the TWSE (there is a net loss during the fiscal year)

The data provided by the TWSE include daily order data, transaction data and disclosed price data from February 2008 to May 2008 (77 trading days). Based on Equation 9.1 and 9.2, we calculated the liquidity costs of all effective market order transactions for the 10 securities. We neglect the effective market orders that arrive before the first best price is disclosed, since the opposing best prices of these orders are not available. The same applies to the effective market orders that arrive after the market is closed. These orders would be recorded on the order book and executed the following day (or later) if their prices have a match.

We are particularly interested in the liquidity costs of effective market orders that were traded immediately right after the orders were submitted, as they were a strong indication of the market liquidity of a security. The liquidity costs for orders, whose transactions took place after the orders have been entered (and waited for) in the order book, might be influenced by other factors, such as opportunity cost, and hence are less informative about the liquidity of a security.

Table 9.7 gives the liquidity cost statistics of the effective market orders with immediate transactions that were executed during the 77 trading days. It shows that most securities have negative average liquidity costs, except the two securities that have the lowest transaction frequency (0100T and 01008T).³ The maximum amount of liquidity cost for a transaction is less than 3%. This indicates that these securities have high market liquidity.

We noted that China Steel (2002.TW) and TSMC (2330.TW) have the highest trading frequency (472,354 & 483,565). This might be due to the fact that 2002.TW is a blue chip stock while 2330.TW has a high market capitalization. They are attractive to domestic and foreign investors who seek secure and stable returns.

³ The results of negative liquidity costs may be a little intriguing, and need to be addressed in further study.

Table 9.7 Descriptive statistics of the liquidity costs of effective market orders with immediate transactions, based on the TWSE data

Ticker	Max	Min	Mean	Std. Dev.	Kurtosis	Sum Sq. Dev	No. of Transactions
0050	0.69%	-0.92%	-0.02%	0.000639	23.17	0.04	101419
0056	2.22%	-1.4%	-0.01%	0.000743	122.77	0.01	16040
01007T	2.55%	-1.58%	0%	0.00091	252.81	0	4912
01008T	1.12%	-1.68%	0.01%	0.001133	90.85	0	872
2002	1.7%	-1.89%	-0.05%	0.001119	52.72	0.59	472354
2330	1.73%	-1.88%	-0.06%	0.001372	48.72	0.91	483565
2454	1.33%	-1.47%	-0.04%	0.001525	20.11	0.95	409985
2498	1.5%	-1.63%	-0.04%	0.001297	26.46	0.4	237582
2912	2.21%	-2.67%	-0.07%	0.00239	24.12	0.31	54480
3474	2.13%	-2.97%	-0.05%	0.002128	50.34	0.37	81450

Table 9.8 Ratio of the transaction volume of the effective market orders with immediate transactions to the total transaction volume of all orders, based on the TWSE data

LC range	(-3%,-2%]		(-2%,-1%]		(-1%,0%]		(0%,1%]		(1%,2%]		(2%,3%]	
	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max
0050	-	-	-	-	45.52%	65.88%	6.15%	24.51%	-	-	-	-
0056	-	-	0.01%	0.71%	40.88%	79.82%	9.22%	52.45%	0.01%	1%	0.05%	0.05%
01007T	-	-	0%	0.12%	44.7%	98.79%	4.17%	50.83%	0.01%	2.22%	0.04%	0.04%
01008T	-	-	0.05%	7.35%	45.39%	100%	2.59%	71.43%	0.02%	3.45%	-	-
2002	-	-	0.14%	12.74%	50.54%	71.49%	4.45%	26.33%	0.21%	18.4%	-	-
2330	-	-	0.46%	17.1%	47.38%	69.5%	4.07%	33.37%	0.43%	24.09%	-	-
2454	-	-	0.09%	4.56%	51.34%	68.51%	7.75%	21.25%	0.14%	8.45%	-	-
2498	-	-	0.02%	1.51%	48.62%	67.43%	5.59%	18.9%	0.05%	5.81%	-	-
2912	0.09%	8.9%	0.14%	6.54%	44.18%	77.66%	2.7%	21.82%	0.18%	11.41%	0.01%	0.01%
3474	0.12%	13.68%	0.17%	8.91%	45.54%	67.54%	7.03%	50.54%	0.51%	17.13%	0.35%	0.35%

To further analyze the market liquidity of these 10 securities, we partitioned the liquidity costs into 6 different value ranges. After that, we computed the ratio of the transaction volume of the effective market orders with immediate transactions to the total transaction volume of all orders (which include limit orders, not immediately executed market orders and so on). As shown in Table 9.8, the trading volume of this type of effective market order is more than 50% of the total transaction volume for all 10 securities. Meanwhile, the majority (40-50%) of all transaction volumes for these securities have their liquidity costs between -1% and 0% (see the 6th column of Table 9.8). These statistics further support them as high market liquidity securities.

9.5 Experimental Setup

Based on the definition in Table 9.5, we calculated $p(\mu)$, $p(\alpha)$, $p(\delta)$, $p(n)$ and σ for the 10 securities in Table 9.9. Although TSMC (2330.TW) is the largest market capitalization stock and is traded frequently, the average daily order size is 14,000 shares, which is much lower than that of the Taiwan Top50 Tracker Fund (0050.TW) (45,000 shares). This might be because the Taiwan Top50 Tracker Fund (0050.TW) is mostly traded by market makers,

who normally trade orders with a large volume, while TSMC (2330.TW) is traded by many different kinds of investors. For each of the securities, we made 10 simulation runs, each of which simulates one trading day for the TWSE. The simulation results are presented and analyzed in the following section.

Table 9.9 Computed $p(\alpha)$, $p(\mu)$, $p(\delta)$, $p(n)$ and σ for the 10 studied securities

Ticker	$p(\alpha)$ $\times 10^{-3}$	$p(\mu)$ $\times 10^{-3}$	$p(\delta)$ $\times 10^{-3}$	$p(n)$ $\times 10^{-3}$	σ $\times 10^3$
0050	2.498	1.029	1.586	994.887	45
0056	0.404	0.147	0.232	999.217	19
01007T	0.109	0.059	0.042	999.79	21
01008T	0.026	0.01	0.003	999.961	22
2002	6.292	5.007	1.923	986.778	10
2330	6.676	5.059	2.218	986.047	14
2454	4.401	4.404	2.086	989.109	3
2498	2.57	2.404	1.182	993.844	3
2912	0.9	0.644	0.414	998.042	5
3474	1.195	0.96	0.459	997.386	11

9.6 Simulation Results and Analysis

Using the 10 days of simulation data, we calculated the daily average trading volume and the daily average number of transactions of the effective market orders. We then compared them with that calculated from the 77 days of TWSE data. As shown in Table 9.10, the results calculated from the simulation data are higher than those calculated from the TWSE data for almost all of the securities. This might be because, in our system, a crossing order could be split into effective market orders and effective limit orders (see Section 9.3.1). In most cases, the effective market order volume (shares) is smaller than the volume (shares) of the effective limit order. However, the DFGIS model assumes that the effective market order size is the same as the effective limit order size σ , which is calculated as the average size of both types of orders. In other words, the average order rate for effective market orders is overestimated. Consequently, the effective market order volume simulated based on σ is higher than that for the real TWSE data. One way to address this issue is to use two different order size parameters, one for effective market orders and one for effective limit orders, to conduct our simulation. We will investigate this option in our future work.

We also evaluated the liquidity cost statistics of effective market orders with immediate transactions, based on the simulation data (see Table 9.11). Compared to Table 9.7, the liquidity costs generated by the simulation data are higher than those for the TWSE data. This might also be due to the overestimated effective market order size σ in our system. With a higher effective market order size, the liquidity costs of effective market orders with immediate transactions are likely to be higher.

To rigorously evaluate the similarity between the simulated liquidity costs and the TWSE data, we performed the Mann-Whitney-Wilcoxon (MWW) test on all 10 securities. The resulting p-values are 0 across all 10 securities, indicating they are indeed different from one another.

Table 9.10 Daily effective market order trading volumes and the number of transactions comparison

Ticker	TWSE data		Simulation Data	
	daily trading volume (share)	daily no. of transactions	daily trading volume (share)	daily no. of transactions
0050	15411781	2554	73289200	2997
0056	1438702	352	5334100	506
01007T	1731457	144	2136900	190
01008T	332438	24	447200	34
2002	50360139	12062	64442500	11826
2330	69341433	12582	95335300	12765
2454	12749066	8010	17604900	7809
2498	6475480	4332	10357200	4788
2912	3228883	1369	5120400	1819
3474	11124524	2183	15510000	2794

Table 9.11 Descriptive statistics of the liquidity costs of effective market orders with immediate transactions, based on the simulation data

Ticker	Max	Min	Mean	Std. Dev.	Kurtosis	Sum Sq. Dev	No. of Transactions
0050	2.66%	-3.05%	-0.12%	0.006948	4.43	0.79	16470
0056	4.48%	-4.41%	0.18%	0.008189	9.38	0.16	2359
01007T	9.65%	-6.42%	0.5%	0.01279	14.84	0.15	889
01008T	6.39%	-1.94%	0.74%	0.012946	8.15	0.02	129
2002	1.79%	-2.22%	-0.18%	0.00562	3.42	2.56	81047
2330	2.03%	-2.35%	-0.21%	0.006375	3.37	3.31	81490
2454	2.93%	-3.62%	-0.28%	0.008793	3.45	5.51	71213
2498	3.76%	-4.65%	-0.3%	0.011251	3.78	4.9	38727
2912	7.32%	-11.11%	-0.11%	0.015314	8.4	2.42	10320
3474	5.35%	-6.43%	-0.28%	0.015003	4.38	3.5	15556

Similar to Table 9.8, we computed the ratio of the transaction volume of the effective market orders with immediate transactions to the total transaction volume for all orders under liquidity cost for 6 different ranges. However, the value ranges are partitioned slightly differently from those of Table 9.8. This is because the liquidity costs of the simulation data have a wider spread (-11.11%, 9.65%) than those of the TWSE data. Since we are more interested in positive liquidity costs, which are indicators of poor market liquidity, we grouped the negative liquidity costs into one bin and added two bins for liquidity costs beyond 3%. The results are given in Table 9.12.

As shown, the liquidity cost upper bound of 2002 (China Steel) and 2330 (TSMC) is 2%, which is the same for both simulation and TWSE data. Meanwhile, the two securities have higher liquidity cost transactions ratios (1% - 2%) for simulation and TWSE data that are similar to each other (the difference is $\sim 1\%$). Similarly, the negative liquidity cost transaction ratios (12% - 0%) for the simulation and TWSE data these two securities are not too far from

Table 9.12 Ratio of the transaction volume of the effective market orders with immediate transactions to the total transaction volume of all orders, based on the simulation and TWSE data

LC range	(-12%,0%]		(0%,1%]		(1%,2%]		(2%,3%]		(3%,4%]		(4%,10%]	
	Sim	TWSE	Sim	TWSE	Sim	TWSE	Sim	TWSE	Sim	TWSE	Sim	TWSE
0050	26.99%	45.52%	21.39%	6.15%	2.69%	-	0.08%	-	-	-	-	-
0056	13.53%	40.89%	22.7%	9.22%	5.09%	0.01%	1.11%	0.05%	0.28%	-	0.06%	-
01007T	13.51%	44.7%	17.21%	4.17%	7.72%	0.01%	3.28%	0.04%	1.57%	-	1.26%	-
01008T	9.25%	45.44%	13.14%	2.59%	4.42%	0.02%	3.25%	-	2.09%	-	2.94%	-
2002	44.06%	50.68%	20.85%	4.45%	1.21%	0.21%	-	-	-	-	-	-
2330	44.68%	47.84%	19.61%	4.07%	1.96%	0.43%	0.01%	-	-	-	-	-
2454	45.68%	51.43%	20.48%	7.75%	4.54%	0.14%	0.4%	-	-	-	-	-
2498	40.08%	48.64%	19.1%	5.59%	4.99%	0.05%	1.33%	-	0.21%	-	-	-
2912	29.49%	44.41%	18.31%	2.7%	4.31%	0.18%	2.16%	0.01%	1.03%	-	0.74%	-
3474	31.29%	45.84%	17.09%	7.03%	5.41%	0.51%	2.51%	0.35%	0.96%	-	0.28%	-

each other either. However, they have many more transactions with liquidity costs between 0% – 1% for the simulation data than that for the TWSE data. For an investor, whose main concern is to avert high liquidity costs, the agent-based model produces liquidity costs that are considered to be similar to those for the TWSE data. When devising liquidation strategies for these two securities, this model can be used to simulate liquidity costs under different strategies to identify the optimal ones.

What, then, has distinguished these two securities from others? We examined the data statistics in Table 9.7 and found that they have a high number of transactions. This indicates that this agent-based system simulates liquidity costs more accurately for securities with a higher trading frequency.

9.7 Discussions

The simulated liquidity costs have a wider spread and higher values than those for the TWSE data. This might be due to the following reasons:

1. The average effective market order size rate (σ) used to run the simulation was overestimated. The order-size assumption maintained in this study can also be challenged. In financial literature, order size has been found to follow a power law distribution, and limit order size distribution is different from market order size distribution. By taking into account this difference, we may improve our current simulation.⁴
2. During the simulation, the time interval between two order events is random and independent, which is different from that observed in the real financial markets. Frequently, orders are clustered together in a certain number of time periods, and not evenly distributed throughout a day. This might have contributed to the higher liquidity costs in the simulation data.
3. During simulation, the order events were generated randomly, based on the model parameters, without consulting the order book. This is different from the reality, where an investor normally checks the order book of the opposite side to make sure a profitable matching is possible before submitting an order. In other words, although the order distribution in the simulation system is the same as that for the TWSE (we used the TWSE data to estimate the probability of order submissions), the *sequence* of the order submissions in the simulation system is not optimized as is that devised by human traders.

⁴ This direction is pointed out by one of the referees, and is left for future study.




Fig. 9.2 The transaction volume (share) vs. the number of transactions in the TWSE (top) & Simulation (bottom) data

Consequently, the simulated liquidity costs are likely to be higher than those for the TWSE data.

4. In the simulation system, the price of an effective market order is set to be the highest possible bid (buy order) or the lowest possible ask (sell order) allowed by the TWSE to guarantee an immediate transaction. This hardly happens in reality. Normally, a trader would seek a price that generates a transaction, without going to the extreme of the highest possible bid/lowest possible ask. As a result, the simulated liquidity costs are likely to be higher than those for the TWSE data.

To address the first issue, we plan to use two order size parameters, one for the effective market orders rate and one for the effective limit orders, to improve our system. The second issue has been investigated by Engle and Russell [3]. In particular, they devised an Autoregressive Conditional Duration (ACD) model to more realistically simulate the order arrival time, price and volume in a stock market. We are currently integrating this model to our system. This work will be presented in a different publication.

The analysis of items 3 & 4 suggests that traders who employ intelligence (e.g., incorporating order book information) to make trading decisions in a real stock market produced transactions with lower amounts of liquidity costs than that produced by the zero-intelligent agents in our artificial stock market. To simulate the real market behavior, in terms of the liquidity costs, we need to install intelligence (e.g., learning ability) in the artificial agents in our system. We will explore this avenue of research in future work.

One intelligent behavior demonstrated by the TWSE traders is a more profitable liquidation strategy. As shown in Table 9.12, among the daily total trading volume of the Taiwan Top50 Tracker Fund (0050.T), 51.15% are trading volume from the effective market orders with immediate transactions: 26.99% of which paid negative liquidity cost; 21.39% of which

paid liquidity cost between 0 and 1%; 2.69% of which paid liquidity cost of 1-2%; and 0.08% of which paid liquidity cost of 2-3%.

By contrast, the TWSE data show that 51.67% of the daily total trading volume of this security are trading volume from the effective market orders with immediate transactions: 45.52% of which paid negative liquidity cost; and 6.15% of which paid liquidity cost between 0 and 1%. In other words, given the task of liquidating a large block of securities (around 50% of the daily trading volume in this case), the TWSE traders accomplished the task by paying a lower amount of liquidity cost than the cost paid by the zero-intelligence artificial traders. What strategy has delivered such savings?

We analyzed the TWSE 0050.TW transactions data from effective market orders with immediate transactions on March 20, 2008. Figure 9.2 (top) shows that there are many more small-volume transactions than larger-volume ones. In particular, more than 500 transactions are with 5,000 or 10,000 trading shares. This is very different from the simulation data (see the right of Figure 9.2), where the number of small-volume transactions is not dramatically different from that of the large-volume ones (the scale is 20 to 1). This suggests that TWSE traders submitted many smaller-size orders instead of a large-size order to conduct transactions. This strategy has led to a lower amount of liquidity costs. We plan to incorporate this intelligent behavior into the artificial agents in our system.

9.8 Concluding Remarks

The market liquidity of a security plays an important role in financial investment decisions and in the liquidation strategies of the security. As an alternative to the Algorithmic Trading, this study has developed an agent-based model to examine liquidity costs of stocks and securities traded in the Taiwan Stock Market.

For the 10 TWSE stocks and securities that we studied, the model-simulated liquidity costs are higher than those for the TWSE data. We identified four possible factors that contribute to this result:

- The overestimated effective market order size rate.
- The random market order arrival time designed in the DFGIS model.
- The zero-intelligence of the artificial agents in our model.
- The price of the effective market order.

We can continue improving the model by addressing the above-mentioned issues. A model that behaves in a similar way to the TWSE in terms of the liquidity costs can be used to study liquidity costs and to devise liquidation strategies for stocks and securities traded on the TWSE.

Acknowledgements. The authors are grateful to two anonymous referees for their very helpful suggestions. The research support in the form of Taiwan NSC grant no. NSC 98-2410-H-004-045-MY3 is gratefully acknowledged.

References

1. Almgren, R., Chriss, N.: Optimal execution of portfolio transactions. *Journal of Risk* 3(22), 5–39 (2000)
2. Daniels, M.G., Farmer, J.D., Gillemot, L., Iori, G., Smith, E.: Quantitative model of price diffusion and market friction based on trading as a mechanistic random process. *Physical Review Letters* 90(10), 108102-1–108102-4 (2003)
3. Engle, R.F., Russell, J.R.: Autoregressive conditional duration: a new model for irregularly spaced transaction data. *Econometrica* 66(5), 1127–1162 (1998)
4. Farmer, J.D., Patelli, P., Zovko, I.I.: Supplementary material for The Predictive Power of Zero Intelligence in Financial Markets. *Proceedings of National Academy of Sciences (USA)* 102(6), 2254–2259 (2005)
5. Guo, T.X.: An Agent-Based Simulation of Double-auction Markets. Master's thesis of Graduate Department of Computer Science. University of Toronto (2005)
6. Kalin, D., Zagst, R.: Portfolio optimization under liquidity costs. Working Paper, 04-01, Risklab Germany GmbH, Germany (2004)
7. Weisstein, E.W.: Half-normal Distribution (2005), <http://mathworld.wolfram.com/Half-NormalDistribution.html>