Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

HETEROGENEOUS ASSETS MARKET DESIGN

Yu An

Carey Business School
Johns Hopkins University
Baltimore, MD 21202, USA

Zeyu Zheng

Department of Industrial Engineering
and Operations Research (IEOR)
University of California, Berkeley
Berkeley, CA 94720, USA

ABSTRACT

In this paper, we study market design issues for heterogeneous assets. We consider a setting with multiple assets and multiple traders. The new modeling feature is that a social planner can decide to shut down some asset markets so that no traders can trade these assets. We show that closing down some asset markets can increase the total volume of trades by all traders. This is because a smaller number of assets available to trade can reduce coordination failures between different traders, and therefore improves the overall matching efficiency. Using numerical simulation, we study the optimal numbers of markets to open under different scenarios. This has practical implications towards the optimal design of session-based trading protocols in the corporate and municipal bond markets. This problem naturally leads to the use of simulation to select the best system from all possible market design choices. We implement simulation experiments to demonstrate our finding and analyze sensitivities to several market-relevant parameters.

1 INTRODUCTION

Many asset markets feature a large number of heterogeneous assets. For example, during 2006 to 2017, over 100,000 different corporate bonds are traded. (Source: Trade Reporting And Compliance Engine (TRACE) dataset.) From 1998 to 2012, over one million different municipal bonds are traded. (See, (Li and Schürhoff 2019).) Large corporations like General Electrics and J.P. Morgan have thousands of bonds outstanding. (See, (Prager et al. 2013).) By the fourth quarter of 2018, there are \$9.2 trillion of corporate bonds and \$3.8 trillion of municipal bonds outstanding. In recent years, both markets are moving towards higher degrees of electronic trading. (See, for example, (Bank for International Settlements 2016).) There are interests from both academics and industry to search for better designs for these markets. (See "The bond market has 99 solutions, and one problem" by Chris White, ViableMkts, July 2016.) However, the majority of academic research focuses on market designs with a single type of assets. (See, for example, (Duffie and Zhu 2017), (Babus and Parlatore 2018), (Dworczak 2019), among many others.) In this paper, we study market design issues for heterogeneous asset markets. The goal is to fill the gap in the academic literature and provide guidance for industry practice.

We consider a setting with multiple assets and multiple traders. A social planner can decide to shut down some asset markets so that no traders can trade these assets. Each trader has some stochastic asset demands and preferences that are his private information. Once the social planner decided which assets are available to trade, each trader strategically submits orders across available markets in order to maximize his *individual* expected volume of trades. The objective of the social planner is to choose which markets to open in order to maximize the *total* expected volume of trades across all traders.

We show that closing down some asset markets can increase the total volume of trades across all markets. From an individual trader's perspective, he always finds it privately optimal to split the trading demands evenly across all available assets that he prefers. For a given asset, the marginal matching probability of a trader's order is decreasing in the order size. Therefore, the trader equates margin matching probability across all preferred assets by splitting orders evenly. Intuitively, this is because concentrating orders into a single asset makes a trader compete against himself. On the other hand, the social planner may want to concentrate orders into a few asset markets. For example, a buyer may want to trade assets 1 and 2, and a seller may want to trade assets 1 and 3. If all three asset markets are open, then the buyer's order in asset 2 will miss the seller's order in asset 3. However, if only asset 1 is available for trade, then no such coordination failures between the buyer and seller would occur. Of course, the social planner does not know the exact asset preferences of each trader. In this case, however, we still find that shutting down some asset markets can reduce coordination failures between different traders and thus improve the expected total volume of trades.

The decision to be made by the social planner in practice involves a large number of assets and traders. Simulation can play a vital role in evaluating decision performances and selecting the best market design among a large number of choices. The problem identified by this paper on selecting the set of assets markets open to traders can be formulated as a selection of the best or ranking and selection (R&S) problem that has been widely studied in the simulation literature. R&S problems concern the selection of the best among systems whose performances are unknown and technically impossible to be computed in closed form, so that they can only be learned through simulation sampling. One key issue is on the allocation of sampling resources on different systems, so that the probability of correct selection is statistically guaranteed. See (Kim and Nelson 2006) for a review of R&S methods, and also (Rinott 1978),(Kim and Nelson 2001), (Hong and Nelson 2007), (Luo et al. 2015). These papers mostly focus on simulation designs that are robust and adversarial, in the sense that the probability of correct selection is guaranteed even for the least favorable configuration. The specific configurations arising in the heterogeneous assets market design may present special structures that could potentially be exploited to enhance simulation efficiency. For example, the set of systems to select from can be labeled, say, $\{1,2,\cdots,l\}$, in a way that the corresponding expected performance sequence μ_i, $i=1,2,\cdots,l$ demonstrates a concave pattern. Specifically, $(\mu_k-\mu_j)/(k-j) > (\mu_l-\mu_k)/(l-k)$ for $j<k<l$. Even though the values of μ_i's are unknown a prior, this structure may still be exploited to enhance simulation efficiency. Further, the cost of simulation associated with system i can be increasing in i in a super-linear way, or even geometrically. This gives another incentive to exploit structures such as concavity that may be uniquely presented in such market design problems.

Our paper speaks to the benefit of having some social planners to close down some asset markets. In the corporate bond market, this role is played by electronic trading platforms, and this type of shutting down markets is called "session-based trading." For example, Trumid is a bond trading platform startup in New York that is backed by Peter Thiel and George Soros. A main business model of Trumid is session-based trading. In January 2018, Trumid traded over \$2 billion in volume. (See "Less Is More in World of Corporate Bond Trading, Report Says" by Molly Smith, Bloomberg, March 2018.) Session-based trading concentrates trading to a limited set of instruments in a limited time period. The objective is to aggregate liquidity and trading interests across the markets. The practice implication of our paper is to help guide the optimal design of session-based trading.

Our paper is related to a few papers on market design issues with multiple assets. (Asriyan et al. 2017) show how changing post-trade transparency across correlated assets can affect allocative efficiency in a dynamic setting. Transparency is not the concern of our paper. (Wittwer 2017) shows that whether traders can submit contingent orders across different asset markets does not affect equilibrium allocation and efficiency. In our setting, contingent orders across different markets are not allowed.

2 MODEL

2.1 Model Setup

There are N traders. Each trader $n \in \{1,2,\ldots,N\}$ has a trading demand for Q_n of assets. A positive value of Q_n indicates demand for buying and a negative value indicates demand for selling. The trading demands across different agents are independent. Each random variable Q_n is distributed according to some given probability density function (p.d.f.) $f_n(\cdot) : R \to R^+$, where $f_n(\cdot)$ is an even function. By allowing for different f_n for different traders, one can model both large and small traders in our framework.

There are M assets. Each trader n is endowed with some asset preference

$$I_n = \{a(n,1), a(n,2), \ldots, a(n,k_n)\},$$

where k_n is the number of assets preference by trader n and $a(n,i) \in \{1,2,\ldots,M\}$. A trader prefers to trade each asset with some given probability p_n, independently across assets and across traders. Again, by allowing for different p_n for different traders, one can model a group of traders with wide asset preferences (high p_n) and another group with narrow asset preferences (low p_n).

There is one social planner in the model. The social planner does not observe each trader's trading demands Q_n or asset preferences I_n. Before any trading happens, the social planner can choose which asset markets are open for trading. Given the ex-ante symmetry between assets, we can write the social planner's problem as follows without loss of generality. The social planner chooses an integer $l \in \{1,2,\ldots,M\}$, and then open the asset markets $\{1,2,\ldots,l\}$. The social planner's objective is to maximize the total volume of trades across all agents.

The trading protocol for each asset market is as follows. The short side always trades with probability 1 and the long side is rationed pro rata. That is, suppose there are in total quantity X of buy orders and quantity Y of sell orders for an asset. Then if $X \geq Y$, then all quantity Y of sell orders are fulfilled, and only a fraction Y/X is fulfilled for each unit of buy order. Conversely, if $X < Y$, then all quantity X of buy orders are fulfilled, and only a fraction X/Y is fulfilled for each unit of sell order.

The action of each trader n is as follows. Each trader n observes the realization of his trading demand Q_n, the types I_n of assets that he prefers, and the social planner's choices of l. In this case, we denote the set of available assets to trader n as

$$J_n = I_n \cap \{1,2,\ldots,l\}.$$

If $J_n = \emptyset$, then trader n does not submit any orders. Otherwise, trader n chooses to split the trading demand Q_n across all available markets J_n that are preferred by trader n. Specifically, denote the order submitted by trade n to asset market m as $q_{n,m}$. Then each trader n chooses $\{q_{n,1}, q_{n,2}, \ldots, q_{n,M}\}$ such that

$$q_{n,m} = 0, \text{ if } m \notin J_n,$$

and

$$\sum_{m \in J_n} q_{n,m} = Q_n.$$

Additionally, we impose two requirements. First, we require that

$$q_{n,m} Q_n \geq 0, \text{ if } m \in J_n.$$

That is, if a trader has buying demand $Q_n > 0$, then he only splits buy orders across available markets J_n, and does not put in sell orders. Conversely, if a trader has selling demand $Q_n < 0$, then he only splits sell orders across available markets J_n, and does not put in buy orders. Second, we require that

$$|q_{n,m}| \geq C, \text{ or } q_{n,m} = 0,$$

for some given constant $C > 0$. That is, each asset market has some minimum trading size C. A trader cannot put in an order with size less than C. For example, in the U.S. corporate bond markets, institutional investors such as insurance companies and mutual companies usually submit orders with size of at least \$1 million. Trades with size less than \$1 million on average receive much worse pricing. (See, for example, (Hendershott et al. 2017).) Given these constraints, each trader n simultaneously and independently decide order submission strategy $\{q_{n,1}, q_{n,2}, \ldots, q_{n,M}\}$ in order to maximize his individual expected volume of trade.

The equilibrium concept is Perfect Bayesian Equilibrium. The equilibrium consists of the social planner's choice of l, and the order submission strategy $\{q_{n,1}, q_{n,2}, \ldots, q_{n,M}\}$ of each trader n such that

- The social planner maximizes the total expected volume of trade given all traders' order submission strategy.
- Each trader n maximizes his individual expected volume of trades given all other traders' order submission strategy.

We focus on symmetric equilibrium in which each agent evenly splits orders across all available assets subject to the minimum trading size constraint C. Whenever the minimum trading size constraint binds, a trader randomizes to allocate his orders across available markets. Although we focus on symmetric equilibrium, each trader is allowed to make arbitrary deviations in the order submission strategy.

2.2 Model Solution

2.2.1 Order Submission Strategy of Traders

Our first step is to show that the symmetric strategy of each trader is optimal given that all other traders use symmetric strategy. Denote $X_{n,m}$ as the sum of buy orders submitted by all traders other than trader n in asset market m and $Y_{n,m}$ as the sum of sell orders submitted by all traders other than trader n on asset market m. We can then write down the optimization problem of agent n as

$$\begin{aligned}
&\max_{q_{n,1},q_{n,2},\ldots,q_{n,M}} E\left[\sum_{m\in J_n}\left(\mathbf{1}_{\{q_{n,m}+X_{n,m}>Y_{n,m}\}}\frac{q_{n,m}Y_{n,m}}{q_{n,m}+X_{n,m}}+\mathbf{1}_{\{q_{n,m}+X_{n,m}\le Y_{n,m}\}}q_{n,m}\right)\right] \qquad (1)\\
&\text{subject to} \quad \sum_{m\in J_n} q_{n,m} = Q_n,\\
&\qquad\qquad q_{n,m}Q_n \ge 0, \text{ if } m \in J_n,\\
&\qquad\qquad |q_{n,m}| \ge C, \text{ or } q_{n,m} = 0,\\
&\qquad\qquad q_{n,m} = 0, \text{ if } m \notin J_n.
\end{aligned}$$

Given that all other traders use symmetric strategy, we know that for $m \in \{1, 2, \ldots, l\}$, $(X_{n,m}, Y_{n,m})$ are identically distributed. Therefore, in order to show the optimality of symmetric strategy for trader n, it suffices to show that

$$g(q) \equiv E\left[\left(\mathbf{1}_{\{q+X_{n,m}>Y_{n,m}\}}\frac{qY_{n,m}}{q+X_{n,m}}+\mathbf{1}_{\{q+X_{n,m}\le Y_{n,m}\}}q\right)\right] \qquad (2)$$

is a concave function of q. We denote the p.d.f. of $(X_{n,m}, Y_{n,m})$ as $\varphi(x, y)$. We can then rewrite (2) as

$$g(q) = \int_0^\infty \int_0^{x+q} \frac{qy}{q+x} f(x,y)\,dy\,dx + \int_0^\infty \int_{x+q}^\infty q f(x,y)\,dy\,dx.$$

Taking derivatives over q, we have

$$\begin{aligned} g'(q) &= \int_0^\infty q f(x,x+q)\,dx + \int_0^\infty \int_0^{x+q} \frac{xy}{(q+x)^2} f(x,y)\,dy\,dx \\ &\quad - \int_0^\infty q f(x,x+q)\,dx + \int_0^\infty \int_{x+q}^\infty f(x,y)\,dy\,dx \\ &= \int_0^\infty \int_0^{x+q} \frac{xy}{(q+x)^2} f(x,y)\,dy\,dx + \int_0^\infty \int_{x+q}^\infty f(x,y)\,dy\,dx. \end{aligned}$$

Taking another derivatives over q, we have

$$\begin{aligned} g''(q) &= \int_0^\infty \frac{x}{q+x} f(x,x+q)\,dx - 2\int_0^\infty \int_0^{x+q} \frac{xy}{(q+x)^3} f(x,y)\,dy\,dx \\ &\quad - \int_0^\infty f(x,x+q)\,dy\,dx \\ &= -\int_0^\infty \frac{q}{q+x} f(x,x+q)\,dx - 2\int_0^\infty \int_0^{x+q} \frac{xy}{(q+x)^3} f(x,y)\,dy\,dx \le 0. \end{aligned}$$

Therefore, we see that $g(q)$ is a concave function of q. This shows that the symmetric strategy of each trader is optimal given that all other traders use symmetric strategy.

2.2.2 Optimal Number of Markets - A Decision of Social Planner

We now study the optimal number l of markets to be opened by the social planner. For a given choice of l, we denote the total volume of buy orders submitted by all traders in asset market $m \in \{1,2,\dots,l\}$ as $X_m(l)$, and the total volume of sell orders submitted by all traders in asset market $m \in \{1,2,\dots,l\}$ as $Y_m(l)$. The expected total volume of trades across all asset markets is given by

$$\mu_l \equiv E\left[\sum_{m=1}^{l} \min(X_m(l), Y_m(l))\right]. \tag{3}$$

The social planner's objective is to find the optimal $l \in \{1,2,\dots,M\}$ in order to maximize μ_l of (3). The tradeoff for the social planner is as follows. By choosing a smaller l, some traders may not be able to find any preferred asset among all available assets. By choosing a larger l, the social planner can reduce this asset preference problem. However, conditional on having at least one preferred asset among all available assets, traders are also more likely to split orders across more asset markets. This exacerbates the coordination problem between different traders. The optimal l is determined by the tradeoff between asset preferences and coordination failures.

Given the symmetric strategies of traders, $X_m(l)$ and $Y_m(l)$ can be evaluated numerically, but are too complicated to be calculated analytically. Maximizing (3) corresponds to a ranking and selection simulation problem. In the next section, we discuss a simulation algorithm for the implementation. In Section 4, we present numerical experiments for the optimal choice of l by the social planner and study the sensitivity of optimal l to model parameters.

3 SIMULATION DESIGN

The first goal is to select the optimal design, i.e., the optimal number of markets to open. Each $l \in \{1,2,\dots,M\}$ represents a different system design, and selection routines are employed to choose an l such that μ_l is maximized over all choices. Different from adversarial routines in which the correct selection probability is guaranteed for the least favorable configuration, the market design problems usually demonstrate a concave pattern for the sequence of μ_l, $l = 1,2,\ldots,M$. We conjecture that for the general class of problems defined in Section 2 the concave pattern could be validated. Specifically, the expected performance sequence μ_l, $l = 1,2,\cdots,M$ satisfies that $(\mu_k - \mu_j)/(k-j) > (\mu_l - \mu_k)/(l-k)$ for $1 \leq j < k < l \leq M$. We denote $N_l > 0$ as allocated to the simulation for system l. Let X_{lj} denote the j-th independent observation from system l and $\bar{X}_l(r)$ denote the sample mean of the first r observations from system l. Denote the corresponding estimated performance as $\hat{\mu}_l$, omitting the dependence on sample size. The following result can be obtained via concentration inequalities.

Theorem 1 Under the concavity assumption, suppose that the expected performance sequence is distinguishable by $\delta > 0$ such that

$$|\mu_l - \mu_{l+1}| > \delta, \forall l = 1,2,\cdots,M-1,$$

then there exists an $C > 0$ such that

$$P(\hat{\mu}_l < \hat{\mu}_{l+1} \mid \operatorname*{argmax}_k \mu_k \leq l) \leq e^{-C\delta N_l} + e^{-C\delta N_{l+1}}.$$

This result can be used to enhance existing ranking and selection procedures. The following procedure is adapted from (Rinott 1978) and (Kim and Nelson 2007) and shows how this result can be used to implement an early-stopping rule.

1. First fix a confidence level α, an indifference parameter $\delta > 0$, and a sample size $n_0 \geq 2$ for the first stage.
2. Obtain Rinott's constant $h = h(n_0, M, 1-\alpha)$ from the tables in (Wilcox 1984). Obtain n_0 observations $X_{lj}, j = 1,2,\ldots,n_0$, from each system $l = 1,2,\ldots,M$. For $l = 1,2,\ldots,M$, compute

$$S_l^2 = \frac{1}{n_0 - 1}\sum_{j=1}^{n_0}(X_{lj} - \bar{X}_l(n_0))^2,$$

   and let

$$N_l = \max\left\{n_0, \left\lceil \frac{h^2 S_l^2}{\delta^2} \right\rceil\right\}$$

   where $\lceil \cdot \rceil$ indicates rounding up to the next larger integer.
3. If $n_0 \geq \max_l N_l$, then stop the procedure and select the system with the largest $\bar{X}_l(n_0)$, indicating the best. Otherwise, sequentially from $l = 1$ to $l = M$, take $N_l - n_0$ additional observations from each system l for which $N_l > n_0$. If $\bar{X}_l(N_l) > \bar{X}_{l+1}(N_{l+1})$, then stop and select system l, denoted as $\hat{l}$.

This procedure can guarantee correct selection with probability at least

$$1 - \alpha - e^{-C\delta N_{\hat{l}}} - e^{-C\delta N_{\hat{l}+1}}$$

and may save computation efforts for simulating the rest of systems $\hat{l}+1, \hat{l}+2, \cdots, M$. which could be much larger than those of the first $\hat{l}$ systems.

4 SIMULATION EXPERIMENTS

In this section, we use simulation to numerically solve for the optimal number l of markets and study the sensitivity of optimal l to different model parameters.

We first present the baseline case. The parameters are chosen as in Table 1. The "Symmetric Exp(1)" distribution refers to that of a random variable $Z_1 \cdot Z_2$ in which Z_1 follows an exponential distribution with mean parameter 1 and Z_2 is a Bernoulli random variable taking values 1 and -1 with equal probability. (Other sub-Gaussian and heavy-tailed distributions lead to similar results.) In Figure 1, we plot the simulated estimators of total volume μ_l as l varies through 1 to M. As we can see, the optimal l is 5, which maximizes μ_l for $l = 1, 2, \ldots, 30$. We denote the optimal choice of l as l^*. On the one hand, because the asset preferences of individual traders are not observed by the social planner, the social planner wants to open more asset markets. This is because if none of the markets of a trader's preferred assets is open, the trader's demand is completely unfilled. On the other hand, opening more asset markets also has the harmful effect of increasing coordination failure. This is because each trader has an incentive to split their order flows evenly across all open markets of preferred assets, subject to the minimum trading size requirement. This causes the order flows in each asset market to be thin, and leads to coordination failure that forgoes potential trading opportunities. The optimal l^* is determined by the tradeoff of these two forces. In our numerical example, the optimal l^* is strictly less than the total number of assets M. This supports our main result that closing down some asset markets can strictly improve total volume of trades.

We implemented a second example to show the effect of heterogeneous density function $f(\cdot)$ between different traders. Specifically, we keep all parameters in Table 1 except for that the demand distribution for half the traders is symmetric Exp(0.5) while the other half is symmetric Exp(1.5). Figure 2 plots the simulated estimators of total volume μ_l as l varies through 1 to M, covered by corresponding confidence intervals. As we can see, the optimal l is 6, different from the result in Figure 1. However, the concave pattern is similar.

Table 1: Parameters Setting.

Parameters	M	N	p	Q_n	C
Value	30	30	0.5	Symmetric Exp(1)	0.5

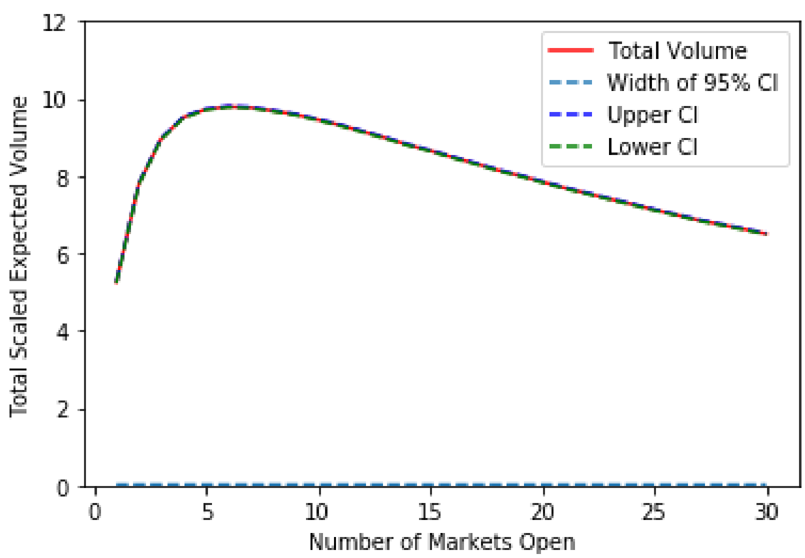

Figure 1: Baseline Case: The "Width of 95% CI" stands for the width of confidence interval at the 95% level. "Upper CI" and "Lower CI" plot the upper and lower range of the associated 95% confidence interval covering the simulated average total volume.

Next, we study the sensitivity of optimal l^* to model papers. We first study the sensitivity to p, which is the probability of any given trader preferring any given asset. Table 2 shows how l^* changes with p. All other parameters are the same as in Table 1. As we can see, higher p leads to smaller l. This is because higher p eases the need to satisfy traders' asset preferences. Given the tradeoff between asset preferences and coordination failure, we see that higher p leads to smaller l.

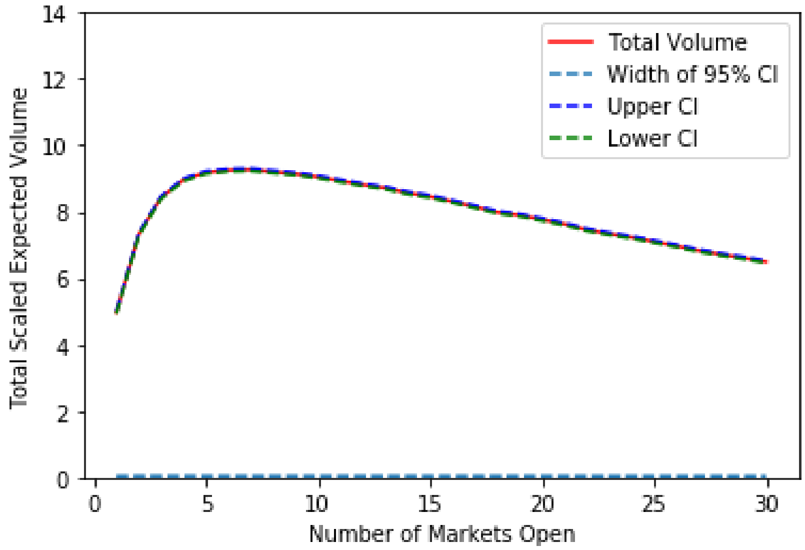

Figure 2: The “Width of 95% CI” stands for the width of confidence interval at the 95% level. “Upper CI” and “Lower CI” plot the upper and lower range of the associated 95% confidence interval covering the simulated average total volume.

Table 2: Sensitivity of Optimal l^* to Parameter p.

Parameter Value	$p = 0.1$	$p = 0.3$	$p = 0.5$	$p = 0.7$	$p = 0.9$
Optimal l^*	26	10	5	3	1

Table 3 shows the sensitivity of l^* to C, the minimal order submission size. As we can see, larger C leads to smaller l. This is because with a larger C, traders need to split their orders into fewer asset markets. This makes it more likely for different traders to miss the trading demands of each other, and therefore exacerbates coordinating failure. As a result, higher C leads to smaller l.

Table 3: Sensitivity of Optimal l^* to Parameter C.

Parameter Value	$C = 0.1$	$C = 0.3$	$C = 0.5$	$C = 0.7$	$C = 0.9$
Optimal l^*	15	9	5	4	3

5 CONCLUSION AND FUTURE WORK

In this paper, we study market design issues for heterogeneous asset markets, in which a social planner decides to open some markets while closing others. We show that closing down some asset markets can increase the total volume of trades. This has practical implications towards the optimal design of session-based trading protocols in the corporate and municipal bond markets. We implement simulation experiments to demonstrate this feature and analyze sensitivity to various model parameters that are relevant to market design.

Our current work could be extended in several directions in the future. First, one could introduce correlation in asset preferences between different traders and different assets, and study how different correlation structures affect the optimal number of markets. This could capture more realistic features of traders and assets. Second, there is no role for prices in our current model for the sake of tractability. One could extend our model to allow for endogenous asset price determination, in order to study the impact on the strategic order submissions by traders and on the optimal number of markets by the social planner. We believe our main results and intuitions would carry through to this more general setting.

REFERENCES

Asriyan, V., W. Fuchs, and B. Green. 2017, July. "Information Spillovers in Asset Markets with Correlated Values". *American Economic Review* 107(7):2007–40.

Babus, A., and C. Parlatore. 2018. "Strategic Fragmented Markets". Working paper, Washington University in St. Louis and New York University.

Bank for International Settlements 2016. "Electronic Trading in Fixed Income Markets". Bank for International Settlements Report.

Duffie, D., and H. Zhu. 2017. "Size Discovery". *The Review of Financial Studies* 30(4):1095–1150.

Dworczak, P. 2019. "Mechanism Design with Aftermarkets: Cutoff Mechanisms". Working paper, Northwestern University.

Hendershott, T., D. Li, D. Livdan, and N. Schürhoff. 2017. "Relationship Trading in OTC Markets". Working paper, University of California Berkeley, Federal Reserve Board and University of Lausanne.

Hong, L. J., and B. L. Nelson. 2007. "Selecting the Best System when Systems are Revealed Sequentially". *IIE Transactions* 39(7):723–734.

Kim, S.-H., and B. L. Nelson. 2001. "A Fully Sequential Procedure for Indifference-Zone Selection in Simulation". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 11(3):251–273.

Kim, S.-H., and B. L. Nelson. 2006. "Selecting the Best System". *Handbooks in operations research and management science* 13:501–534.

Kim, S.-H., and B. L. Nelson. 2007. "Recent Advances in Ranking and Selection". In *Proceedings of the 2007 Winter Simulation Conference*, 162–172. S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, R. R. Barton eds., Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Li, D., and N. Schürhoff. 2019. "Dealer Networks". *The Journal of Finance* 74(1):91–144.

Luo, J., L. J. Hong, B. L. Nelson, and Y. Wu. 2015. "Fully Sequential Procedures for Large-Scale Ranking-and-Selection Problems in Parallel Computing Environments". *Operations Research* 63(5):1177–1194.

Prager, R., K. Riaz, S. Vedbrat, and E. C. Watt. 2013. "Setting New Standards. The Liquidity Challenge II". BlackRock Research Report.

Rinott, Y. 1978. "On Two-stage Selection Procedures and Related Probability-Inequalities". *Communications in Statistics-Theory and methods* 7(8):799–811.

Wilcox, R. R. 1984. "A Table for Rinott's Selection Procedure". *Journal of Quality Technology* 16(2):97–100.

Wittwer, M. 2017. "Centralizing Disconnected Markets? An Irrelevance Result". Working paper, Stanford University.

ACKNOWLEDGEMENT

The authors deeply thank the two anonymous referees and the track chairs for their very helpful suggestions and comments.

AUTHOR BIOGRAPHIES

YU AN is an Assistant Professor of finance at Johns Hopkins Carey Business School. He received his Ph.D in Finance (2019) and M.S. in Financial Mathematics (2014) from Stanford University, and B.A. in Finance (2012) from Peking University. His research interests are in financial economics and over-the-counter markets. His email address is yua@jhu.edu.

ZEYU ZHENG is an Assistant Professor at the Department of Industrial Engineering and Operations Research (IEOR) at the University of California, Berkeley. He received his Ph.D in Management Science

and Engineering (2018), Ph.D. minor in Statistics (2018) and M.S. in Economics (2016) from Stanford University. He received a B.S. in Mathematics (2012) from Peking University. His has research interests in simulation, stochastic modeling, data-oriented decision making, and financial technologies. His email address is zyzheng@berkeley.edu.