

Intermediated Blind Portfolio Auctions

Michael Padilla	Benjamin Van Roy
Stanford University	Stanford University
mtp@stanford.edu	bvr@stanford.edu

Abstract

As much as 12% of the daily volume on the New York Stock Exchange, and similar volumes on other major world exchanges, involves sales by institutional investors to brokers through blind portfolio auctions. Such transactions typically take the form of a first-price sealed-bid auction in which the seller engages a few potential brokers and provides limited information about the portfolio being sold. Uncertainty about the portfolio contents reduces bids, effectively increasing the transaction cost paid by the seller. We consider the use of a trusted intermediary or equivalent cryptographic protocol to reduce transaction costs. In particular, we propose a mechanism through which each party provides relevant private information to an intermediary who ultimately reveals only the portfolio contents and price paid, and only to the seller and winning broker. Through analysis of a game-theoretic model, we demonstrate substantial potential benefits to sellers. For example, under reasonable assumptions a seller can reduce expected transaction costs by more than 10%.

1 Introduction

As the financial world has become increasingly complex and the needs of investors have become more varied, financial markets have responded in kind via the creation of a plethora of alternative liquidity sources. Focusing on the equity markets, in addition to the long-established primary exchanges such as the New York Stock Exchange (NYSE) and the Tokyo

Stock Exchange (TSE), investors' liquidity options are now enriched by the addition of electronic communication networks (ECNs) such as Instinet and NYSE Arca, as well as over the counter (OTC) markets, block/upstairs markets, and the newer addition of "dark pools" (Sofianos (2007)). In addition, the degree of sophistication in the services offered by the finance industry to investors has significantly increased as well, one example being the development of the program¹ trading business.

As institutional investors shift the contents of their portfolios over the course of their investment strategy, they frequently desire to trade a collection of names as an aggregate portfolio rather than individually. While investors typically have the option of executing their portfolio trades via their own discretion and control, typically using a trading tool provided by a broker that allows them to access that broker's link to various exchanges, they will often elect to trade through a broker's program trading service. Potential benefits of this approach arise from leveraging the broker's proprietary technology and trading expertise as well as the avoidance of a variety of execution, price, and tracking risks by transferring them to the broker, as we now describe.

The program trading business operates primarily in two formats, agency and principal. In an *agency* trade, the facilitating broker executes the seller's portfolio as per their instructions for a fixed commission². Any and all risk is held by the seller, as the broker merely executes trades on behalf of the seller and receives a commission. A common instruction from the seller to the broker in this scenario is to try to meet or beat the VWAP³ price of that portfolio over the day. The other primary flavor of program trade is a *principal* trade. This differs markedly from the agency trade in that the seller sells its portfolio to the broker at a particular mutually agreed upon spot value plus a commission, transferring ownership of the portfolio and all inherent risks to the broker. Spot value examples are the previous day's close or the current day's close. The broker's commission is typically quoted as a fee per share.

¹In this context "program" is used synonymously with "portfolio".

²Without loss of generality, we assume the investor is selling the portfolio for simplicity in our discussion.

³Volume Weighted Average Price is the ratio of the value traded to the total volume of shares traded over a particular time horizon (usually one day). It is a measure of the average price a stock (or portfolio) traded at over the trading horizon.

The seller's motivation for participation in a principal trade is to reduce the transaction costs incurred in trading their portfolio, the liquidity of which is a primary determinant. Sellers need to carefully consider the benefits of executing their trades themselves or via an agency trade with a broker, or alternatively by making use of a broker's principal trading option. If liquidity is a concern, due to a combination of trades that compose a large fraction of a day's volume and/or a necessity to execute the portfolio quickly, and they would like to avoid stretching the trade over a longer period and be exposed to the associated risks, then locking in a particular spot value via a principal trade can be particularly attractive. From the broker's point of view, principal trades present an opportunity to not only obtain a commission, but possibly more importantly to gain access to new flows from which they can unwind existing positions or facilitate new business. The key to a broker's success in a principal trade is being able to offer a bid that is both cheaper than what the seller estimates their execution costs to be if they traded the portfolio themselves or via an agency trade, while ensuring that the commission adequately covers the risk of taking possession of the portfolio. The degree to which the broker has pre-existing advantageous positions, a flexible longer-term trading horizon, and superior trading technology and know-how will determine their success.

Such principal trades are often the result of a preliminary *blind portfolio auction*. In a blind portfolio auction, the seller will contact a handful of brokers and provide them with only aggregate features of the portfolio and not its precise contents. Features may include, but are not limited to, the number of names, the size of the portfolio, sector exposures, liquidity, bid-ask spreads, expected transaction costs, tracking error with respect to some index, Barra⁴ exposures, portfolio beta, etc. In some cases, the seller will provide the same information to each broker, while more generally brokers will offer a "pre-trade" tool to sellers which calculates some of the metrics according to their own specific evaluation criteria. After each broker provides its bid to the seller, the seller decides whether to accept the best bid. If accepted, the entire portfolio is sold to the winner. Throughout the process, each broker is unaware of the identities of their competition and the competing bids. Blind portfolio

⁴<http://www.msclub.com/>

auctions account for roughly 12%⁵ of the daily traded volume in shares traded on the NYSE, equivalent to approximately 18 billion dollars daily.

The motivation for this “blind” procedure is to furnish brokers with enough information to calculate a bid price that adequately reflects the inherent risks of the portfolio while not revealing information that would enable them to determine its specific contents. The reasons for this privacy concern are many. One fear is that if a future spot price is used, a broker with sufficient knowledge of the portfolio can *pre-hedge* by trading all or a portion of the portfolio before the auction is even complete. The result of such activity is two-fold: while the broker is able to provide a very competitive and likely winning bid to the seller, pre-hedging will have likely caused market impact, resulting in a worse spot price for the seller. A related concern is that brokers who do not win the auction (or perhaps do not even intend to) may use knowledge of the portfolio’s contents to *front-run* the anticipated portfolio trade⁶. This creates additional risk for each broker should they win and consequently reduces bids, again to the detriment of the seller. It is important to note, however, that even in a blind auction brokers are often able to estimate the actual names, directions, and quantities to varying degrees of accuracy for a subset of the portfolio based on factors such as previous experience with that investor and an understanding of their trading patterns, rumors in the market, distinguishing characteristics in the aggregate description, etc. This inadvertent revelation of portfolio information will be referred to as “information leakage”. Although no paper exploring information leakage in blind portfolio auctions has yet been published, considerable anecdotal evidence from market practitioners as well as the analogous demonstration of information leakage in upstairs markets for large-block transactions presented in Keim and Madhavan (1996) strongly suggests that it is a market reality.

In its current form, these auctions leave room in several dimensions for reduction of the seller’s transaction cost:

⁵Estimated using the NYSE’s program trading reports for January and February 2010. These show the total volume of program trades, broken down into both principal and agency trades. We make the assumption, supported through discussion with industry insiders, that the majority of principal program trades are blind risk trades.

⁶Pre-hedging and front-running are essentially the same activity, trading in front of the pending transaction, but differ in the motivation behind it.

1. **The valuations of risk-averse brokers are reduced due to uncertainty about the portfolio's contents.** Due to the blind nature of the bidding process, brokers are concerned about adverse selection effects. This concern is naturally factored into their bids, resulting in a worse price for the seller. Ideally, we would like to implement a mechanism in which the *bid formation process* of each broker is maximally informed, while each *broker themselves* is maximally uninformed until necessary (i.e. they win the auction).
2. **It is desirable to solicit bids from a larger number of brokers.** Due to the information leakage that occurs in the current mechanism, sellers want to limit their bid solicitation to a limited number of brokers. This number is determined in an ad-hoc manner by the seller by considering the trade-off of being able to choose a winning bid from among a larger pool and, on the other hand, having more brokers privy to information about the pending trade. If an auction mechanism can keep brokers maximally uninformed as to the portfolio's contents during the bidding process then the seller is able to solicit a much larger number of bids, to their advantage.
3. **The value of the portfolio may be increased through division into parts sold to multiple brokers.** Currently, the seller awards the entire portfolio to the winning broker. The winning broker's bid is based on a variety of factors, including their proprietary and seller-based flows and ability to hedge and cross with such flows, their own idiosyncratic risks and trading strategies, etc. In the case of a large portfolio, it is likely that different pieces of the portfolio are valued more by different brokers. As a result, the seller can potentially receive a better overall bid for their portfolio by allocating different parts to different brokers. This may be conceptually viewed in the framework of a package auction (Milgrom (2004)), in which an auction mechanism is designed to allocate the objects (e.g., portfolio components) among the brokers so as to maximize the overall social welfare of the participants, which will translate to acquiring the best overall bid for the seller.

This paper studies how factors 1 and 2 might be addressed when the auction makes use of an intermediary. The notion of a centralized intermediary orchestrating the auction is

facilitated by assuming that each broker's bid formation process is highly quantitative and can be encapsulated into a function for which the input is the portfolio being auctioned and the output is the broker's bid. An interesting real-world example of this is a method known to have been practiced by the brokerage D.E. Shaw, in which a computer program was provided to the seller. Upon the seller entering their (actual) portfolio into the program it would produce an encrypted number and aggregate report which would then be sent to D.E. Shaw, who would then decrypt the number to produce the bid value, presumably checked via inspection of the aggregate report (Traders Magazine (1998)). As brokers become increasingly quantitative in their processes, the ability to encapsulate their bid calculation procedure into a program of this sort will become increasingly natural and commonplace. Note that in such scenarios brokers would not want to hand their programs to the seller and allow them to produce an unencrypted bid due to concerns that this could allow the seller to potentially use it to reverse-engineer proprietary information.

With this assumption that brokers are able to represent their bid formation processes as stand-alone programs (or equivalently as mathematical functions) that may be handed to a trusted intermediary, this paper proposes and studies a blind auction format in which the bid formation process of each broker is maximally informed, while each broker themselves is maximally uninformed, denied explicit knowledge of the portfolio until necessary (i.e., when they win the auction). This realizes the goal of eliminating any information leakage that occurs in the current auction process. Throughout the bidding process, both the seller and brokers are left completely uninformed of each others' positions and bids, with the final result that only the seller and winning broker (if any; the seller may reject all bids) know the contents of the seller's portfolio and the winning bid. The mechanism may be implemented in either one of two ways: via a trusted intermediary that stands between brokers and the seller, or by an equivalent cryptographic construct.

It is not surprising that a trusted intermediary can facilitate a reduction in the seller's transaction cost by addressing the factors listed above. However, it is not entirely clear whether the magnitude of this benefit is sufficient to warrant significant attention from industry participants and, in particular, the major effort required to restructure how such transactions are conducted. Our analysis in this paper demonstrates that the potential

benefits can indeed be large. We find that under reasonable conditions, if there are two participating brokers and the seller must sell his entire portfolio to one or the other, transaction costs can be reduced by as much as 10%.

Our study is computational, and a contribution of this paper is in formulating a simple model that captures essential features of the problem under consideration and the development of an approach that efficiently computes Bayesian-Nash equilibrium for auctions with and without an intermediary. Aside from quantifying potential transaction cost reductions for a fixed number of brokers, our computational results shed light on the impact of the number of brokers and of risk aversion. As one would expect, with or without an intermediary, transaction costs diminish greatly as the number of participating brokers increases; going from two to eight brokers reduces transaction costs by about a factor of three. This suggests major additional advantages to using an intermediary if that makes a seller more comfortable to engage a larger number of brokers.

As for the impact of risk aversion, our computational results point to a striking difference between auctions with and without an intermediary. In the absence of an intermediary, seller transaction costs increase with broker risk aversion. On the other hand, with an intermediary, seller transaction costs decrease with broker risk aversion. This somewhat surprising fact gives rise to additional advantages from use of an intermediary. We prove a theorem establishing that this result holds in a fairly general setting.

This paper is organized as follows. In Section 2 we provide an overview of prior work on blind portfolio auctions. Section 3 then provides a case study to motivate the need for informational privacy in our auction format. The model used in our study is described in Section 4 and results characterizing equilibrium bidding strategies are presented in Section 5. Section 6 then presents and discusses results. Two approaches by which an intermediary can be implemented are discussed in Section 7, followed by our conclusion in Section 8.

2 Literature Review

In spite of the fact that blind portfolio trades account for approximately 12% of the daily shares traded volume on the NYSE, this significant business practice has received little

attention in the academic literature. This is largely due to the highly proprietary nature of these transactions and the resulting paucity of publicly available data, as both buy and sell-side participants prefer to retain any competitive advantage possible by maintaining data confidentiality. That said, there is a slowly growing community of researchers looking at this important area of finance. This paper is to our knowledge the first investigation of how the notion of uncertainty affects the performance of blind portfolio auctions and to consider alternative implementations to ameliorate these effects.

In one of the earlier papers on this topic, Almgren and Chriss (2003) developed a mathematical framework for pricing and trading principle portfolios, annualizing the performance of the potential trade and placing it into the context of a firm evaluating this trading opportunity alongside other possible investments. The primary pricing tool developed is that of the information ratio, which measures the balance of the expected value versus the variance of the transaction cost of trading a given portfolio, akin to the well-known Sharpe Ratio⁷. Analogous to Markowitzian portfolio theory, they introduce the notion of an efficient trading frontier, parameterized by the time until liquidation of the portfolio. The information ratio provides a measure of risk-adjusted expected profitability and may be used by an investor when deciding to engage in a particular trade.

As the first empirical study, Kavajecz and Keim (2005) investigated the transaction cost savings of using blind portfolio trades as opposed to more traditional single-name trading methods and analyzed the sources of the 48% reduction in costs they found were realized on average. This study models the cost of trading single-names via an empirical regression model based on portfolio characteristics and also uses actual data for 83 trades over a two year period received from a single asset manager. They found evidence that lower transaction costs are made possible by brokers' ability to leverage several advantages, namely the ability to cross shares with internal inventory, having a longer time horizon to trade out of the position, and superior trading ability.

Continuing this line of empirical study, Giannikos and Suen (2007a), Giannikos and Suen (2007b), and Giannikos et al. (2009) investigate how the empirical model in Kavajecz and Keim (2005) as well as the structural models developed to explain a market maker's spread

⁷The Sharpe Ratio is the expected return divided by its standard deviation.

in Stoll (1978a), Stoll (1978b), and Bollen et al. (2004) can be applied with modification to also model brokers' bids. The authors make an interesting observation from studying the analogy between the role of a market maker and that of a broker bidding on a blind portfolio; in both cases they are facilitating a counterparty's need for immediacy and in return need to be compensated for both inventory and adverse selection risks. In developing their empirical model, they go beyond Kavajecz and Keim (2005) by including additional determinants that account for factors in the overall trading environment, not just the auctioned portfolio itself, that would also affect a broker's risk exposure. In their structural model based on Bollen et al. (2004), they are able to analyze the bid into various fundamental and structural cost components, the largest of which is found to be the inventory holding costs. In addition, Suen (2007) studies a portfolio manager's choice in deciding between an agency or blind principle trade, finding empirical evidence that prospect theory better models the decision making process than does expected utility theory.

3 Blind Auction Example

To develop an understanding of some of the concerns in a blind portfolio auction, consider the example of an actual (buy) trade of a portfolio of 55 FTSE 100 stocks discussed in Ferry (2004) and shown graphically in Figure 1. In this case, a broker is contacted in the morning about a portfolio that is to be guaranteed a price marked to a specific time in the early afternoon. The broker is given partial information about the portfolio being traded and from this is able to determine both the direction (buy) as well as the actual names for a subset of the stocks in the portfolio (information leakage). Confident in their estimate of this subset of the portfolio, presumably composing a large fraction of the portfolio's volume, they begin to buy those stocks ahead of the portfolio being awarded to any of the brokers. This allows the broker the chance to buy the shares necessary for the trade earlier and more gradually relative to their competition in preparation for making an aggressive bid that hopefully wins, i.e. pre-hedging. In effect, the broker transfers all or some of the associated market impact of the trade to the buyer by purchasing shares early and affecting the spot price. This has a deleterious effect for the investor as the price that he pays for his buy

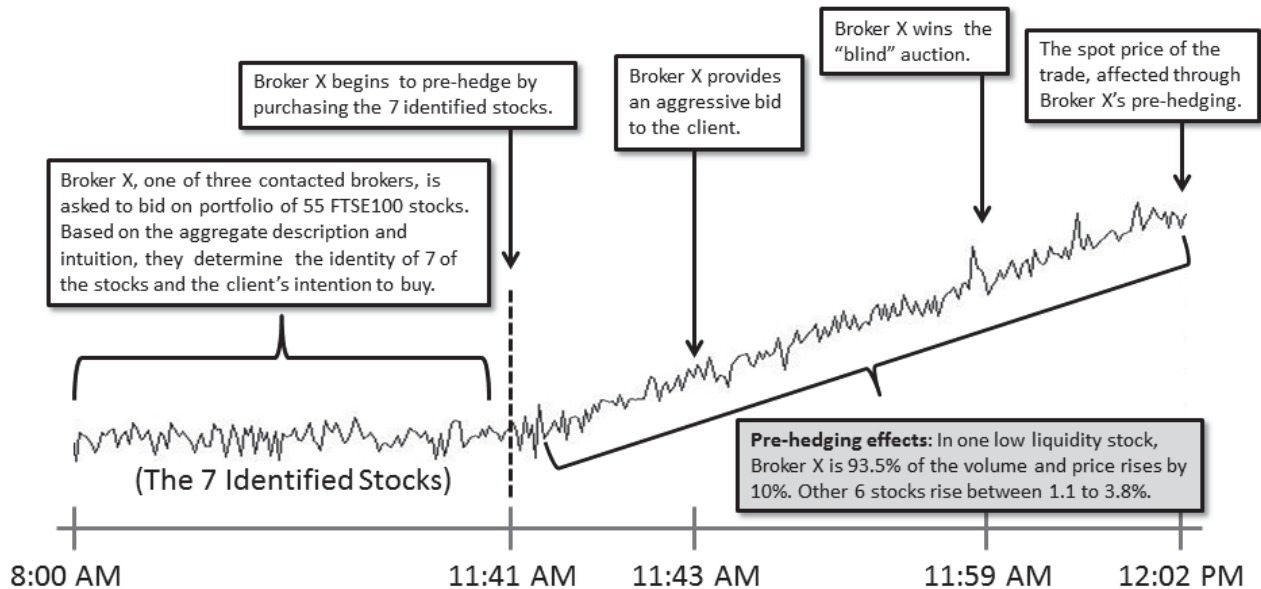


Figure 1: *Representative example of the effects of pre-hedging (Ferry (2004)).*

portfolio has now been “artificially” increased due to the pre-hedging.

A similar effect occurs when non-winning brokers are able to identify the names and direction of some portion of the portfolio and begin to buy ahead of the portfolio being awarded to another broker, i.e. front-running. The purpose is not to cheaply buy shares that will later be exchanged with the investor, but rather to buy shares cheaply and then later sell them to the winning broker who will need to buy the shares necessary to do the trade with the investor. This second form of pre-trading is of course a concern not only to the investor, but to the winning broker as well as it represents price-risk that they absorb upon winning the auction. As such, this risk is factored into brokers’ bids.

4 Model

Our goal is to demonstrate the value that an intermediary can introduce by eliminating the impact of uncertainty about portfolio contents. In this section, we formulate a simple model that captures the essential characteristics of the problem in a way that is sufficient for this purpose.

We consider a situation with a single seller and N brokers, each of whom bids on the

seller's portfolio. The seller is risk-neutral⁸ and brokers share identical preferences represented by a utility function $u : \mathfrak{R} \mapsto \mathfrak{R}$. To facilitate our analysis, some of our results will be established under the following assumption.

Assumption 4.1. *Brokers exhibit constant absolute risk aversion. In other words, there exists a scalar $r > 0$ such that $u(v) = -\exp(-rv)$ for all v .*

Let v^* denote the nominal value of the portfolio. We take r and v^* to be fixed constants, known to the seller and all brokers. In the standard auction without an intermediary, we assume both for convenience and to focus our study on the uncertainty regarding portfolio characteristics that N is known by the brokers. In the intermediated auction, brokers are not assumed to know N as it is instead supplied to brokers' bidding functions via proxy by the intermediary as is explained later in Section 5.2. The role of private information is central to our problem and arises through the dollar value v_n of the portfolio to each n th broker. We will make the following simplifying assumption about the information structure.

Assumption 4.2. *For each n , the value of the portfolio to the n th broker is given by $v_n = v^* - q\theta_n$. The random variables $\theta_1, \dots, \theta_N$ are iid with bounded support, denoted by Θ , and are independent of the random variable q , which we assume to also have bounded support. Both are positive with probability one and have non-zero measure. Further, the seller only observes q and each n th broker only observes θ_n .*

The parameter q captures characteristics of the portfolio driving its desirability to brokers, and θ_n is a parameter that captures the sensitivity of broker n to q . In reality of course a given portfolio may have associated with it a vector of relevant descriptors $\bar{q} = [q_1, \dots, q_D]$, where $q_d \in \mathfrak{R}$ and D is the number of descriptors, and we might think of a given broker's valuation as a function mapping $\mathfrak{R}^D \mapsto \mathfrak{R}$. If we make the assumption that the broker's valuation function is of the form $v(\bar{q}) = \sum_{d=1}^D \psi_d(q_d)$ where ψ_d is a decreasing linear function for the d th descriptor, then we can equivalently represent \bar{q} as a single descriptor $q \in \mathfrak{R}$ which describes the extent that \bar{q} lies along the direction of greatest decrease of v . The reduction of $\bar{q} \in \mathfrak{R}^D$ to $q \in \mathfrak{R}$ thus amounts to a coordinate transformation. It is this concept that we

⁸This assumption is made for expositional simplicity but none of the results discussed in this paper depend on it as all mechanism dynamics are driven by broker preferences.

are trying to compactly represent by using a single aggregate portfolio descriptor q . Since θ_n is the only payoff-relevant information that is unique to the n th broker, this parameter can be thought of as the broker's type.

Let us offer a more concrete interpretation of the terms defining v_n . Think of v^* as the spot value of the portfolio, which is the commonly known market value that the seller could receive if he could liquidate without incurring transaction costs. The value v_n of the portfolio to a broker is generally less than this. We take v^* to be a fixed constant known to all participants, true for an earlier spot but an allowable simplifying assumption for the case of a future spot since any uncertainty in v^* would be present in either auction mechanism. This assumption is easily removed if desired in our model and can be thought of as reflecting the trend of sellers increasingly preferring past spots. The parameter q , representing portfolio characteristics that determine its value to brokers, can for example be thought of as representing the liquidity of assets in the portfolio. Note that this particular interpretation is for illustration and neither necessary nor all-inclusive. If $q = 0$, the assets are perfectly liquid and $v_n = v^*$. The sensitivity parameter θ_n in turn reflects how rapidly the broker's assessment of the portfolio's value decays with increasing illiquidity. It may depend for example on features of the broker's execution technology and trading expertise. Under this interpretation, the term $\theta_n q$ hence represents the transaction costs incurred by the broker from the time the portfolio is acquired until it is completely liquidated, i.e. the difference in realized value when the broker trades out of the portfolio relative to its spot value v^* when received from the seller.

It is useful to think about the value of the portfolio to broker n prior to observing q in terms of the certainty equivalent under Assumption 4.1:1

$$v_{ce}(\theta_n) = -\frac{1}{r} \ln \left(-E \left[-e^{-rv_n} \mid \theta_n \right] \right). \quad (1)$$

Note that, given only knowledge of θ_n , the n th broker is indifferent between receiving the portfolio and receiving a fixed sum of money $v_{ce}(\theta_n)$. Further, if the broker pays an amount b for the portfolio, the certainty equivalent value of the transaction is

$$-\frac{1}{r} \ln \left(-E \left[-e^{-r(v_n-b)} \mid \theta_n \right] \right) = v_{ce}(\theta_n) - b.$$

5 Mechanisms

In this section we present two auction formats for which we will derive pure strategy bidding functions for a Bayesian-Nash equilibrium, the assumed solution concept we use to predict broker behavior. We begin with the standard first-price auction and then describe its counterpart that makes use of a trusted intermediary.

5.1 Standard First-Price Auction

In the first-price auction, each broker places a bid and the portfolio is sold to the highest broker at the highest bid price. Though this does not affect our analysis, it is worth noting that this may be thought of as a sealed bid auction as each broker generally would not want others to see his bid.

As a model of broker behavior, we will assume brokers behave according to pure strategy bidding functions that form part of a Bayesian-Nash equilibrium. Specifically, given N , r , and v^* and the symmetric nature of the game, we will assume that brokers employ a symmetric strategy which takes the form of a function β_S that maps the private information θ_n of each n th broker to a bid $\beta_S(\theta_n)$, $\beta_S : \mathfrak{R} \mapsto \mathfrak{R}$. The subscript here indicates association with the standard first-price auction. Note that in this section it is assumed that brokers know the value of N as discussed in Section 4.

While in general one might hypothesize the existence of multiple symmetric equilibria as well as possible asymmetric equilibria, we will argue invoking prior results in the literature that there exists a unique Bayesian-Nash equilibrium bidding function, which is symmetric, for this auction. As such, this bidding function β_S uniquely satisfies

$$\beta_S(\theta_n) \in \operatorname{argmax}_{b \in \mathfrak{R}} E \left[u \left((v_n - b) \mathbf{1}(b \geq \max_{m \neq n} \beta_S(\theta_m)) \right) \mid \theta_n \right],$$

for all θ_n . We begin with the following preliminary lemma.

Lemma 5.1. *Let Assumptions 4.1 and 4.2 hold. There exists a unique pure strategy symmetric Bayesian-Nash equilibrium bidding function for the standard first-price auction. This bidding function is strictly decreasing on Θ and is differentiable.*

Proof. This follows from Theorem 2 in Maskin and Riley (1984), the conditions for which our model may be shown to meet. \square

We now derive an expression that characterizes this equilibrium bidding function and will facilitate our computational study.

Theorem 5.2. (Brokers' Bidding Function in a Standard First-Price Auction) *Let Assumptions 4.1 and 4.2 hold and assume that the distributions of broker types θ_n and the portfolio characteristic q have densities. In a standard first-price auction, the unique pure strategy symmetric Bayesian-Nash equilibrium bidding function is given by*

$$\beta_S(\theta_n) = v_{ce}(\theta_n) + \frac{1}{r} \ln \left[1 - \frac{r e^{-r v_{ce}(\theta_n)}}{F_{ce}(v_{ce}(\theta_n))^{N-1}} \int_0^{v_{ce}(\theta_n)} F_{ce}(\rho)^{N-1} e^{r\rho} d\rho \right], \quad (2)$$

where F_{ce} is the cumulative distribution function of each broker's certainty equivalent value.

Proof. The symmetric equilibrium bidding function β satisfies

$$\begin{aligned} \beta(\theta_n) &\in \operatorname{argmax}_{b \in \mathfrak{R}} E \left[-e^{-r(v_n - b)} \mathbf{1}(b \geq \max_{m \neq n} \beta(\theta_m)) \mid \theta_n \right] \\ &\in \operatorname{argmax}_{b \in \mathfrak{R}} P(b \geq \max_{m \neq n} \beta(\theta_m)) \left(E \left[-e^{-r(v_n - b)} \mid \theta_n \right] + 1 \right) - 1, \end{aligned}$$

for all θ_n , where we temporarily drop the subscript S to simplify notation. By Lemma 5.1, there is a unique solution β to these inclusions, and it is monotonically decreasing on Θ and differentiable. Let $\tilde{\beta}$ denote the image of Θ generated by β .

Note that there is an injective mapping between θ_n and values of $v_{ce}^n \triangleq v_{ce}(\theta_n)$ which allows us to reinterpret β equivalently as a mapping from brokers' certainty equivalent valuations to bids. Let \tilde{V}_{ce} denote the image of Θ under the transformation to certainty equivalent value. We now view β with this interpretation and will first solve for the bidding function in terms of v_{ce}^n for convenience, and then as a final step rephrase this in terms of θ_n via substitution. An equivalent condition on the Bayesian-Nash equilibrium bidding function is thus given by

$$\beta(v_{ce}^n) \in \operatorname{argmax}_{b \in \mathfrak{R}} P(b \geq \max_{m \neq n} \beta(v_{ce}^m)) \left(-e^{-r(v_{ce}^n - b)} + 1 \right) - 1.$$

Since the mapping from θ_n to v_{ce}^n is injective, Lemma 5.1 also applies to this reinterpretation of β as a function of v_{ce}^n , providing strictly increasing monotonicity in v_{ce}^n and differentiability.

Since β is monotonically increasing in v_{ce}^n , it has an inverse which we will denote by α . Note that α is differentiable on $\tilde{\beta}$.

A broker who bids b and assumes the other $N - 1$ brokers employ the equilibrium bid function β believes that he will win the auction with probability $F_{CE}(\alpha(b))^{N-1}$, i.e. the probability that all other brokers' certainty equivalent valuations lead to bids lower than his. We therefore have

$$P(b \geq \max_{m \neq n} \beta(v_{ce}^m)) (-e^{-r(v_{ce}^n - b)} + 1) - 1 = F_{CE}(\alpha(b))^{N-1} (-e^{-r(v_{ce}^n - b)} + 1) - 1.$$

If $b = \beta(v_{ce}^n)$ for some $v_{ce}^n \in \tilde{V}_{ce}$, it must satisfy the first-order condition for Bayesian-Nash equilibrium:

$$\begin{aligned} 0 &= \frac{d}{db} \{ F_{CE}(\alpha(b))^{N-1} [-e^{-r(v_{ce}^n - b)} + 1] - 1 \} \\ &= (N - 1) F_{CE}(\alpha(b))^{N-2} \frac{d}{db} F_{ce}(\alpha(b)) [-e^{-r(v_{ce}^n - b)} + 1] + F_{CE}(\alpha(b))^{N-1} [-r e^{-r(v_{ce}^n - b)}] \\ &= (N - 1) F_{CE}(\alpha(b))^{N-2} f_{ce}(\alpha(b)) \alpha'(b) [-e^{-r(v_{ce}^n - b)} + 1] - r F_{CE}(\alpha(b))^{N-1} e^{-r(v_{ce}^n - b)} \end{aligned}$$

Substituting $v_{ce}^n \rightarrow v_{ce}$, $\alpha(b) \rightarrow v_{ce}$, $b \rightarrow \beta(v_{ce})$, $\alpha'(b) \rightarrow 1/\beta'(v_{ce})$, and $\frac{d}{db} F(\alpha(b)) \rightarrow f(\alpha(b))\alpha'(b)$ in the above first-order condition and rearranging terms, we arrive at

$$(N - 1) F_{CE}(v_{ce})^{N-2} f_{ce}(v_{ce}) [-e^{-r(v_{ce} - \beta(v_{ce}))} + 1] - r \beta'(v_{ce}) F_{CE}(v_{ce})^{N-1} e^{-r(v_{ce} - \beta(v_{ce}))} = 0.$$

Let $\zeta(v_{ce}) = e^{r\beta(v_{ce})} F_{ce}(v_{ce})^{N-1}$ so that

$$\zeta'(\theta) = r F_{ce}(v_{ce})^{N-1} \beta'(v_{ce}) e^{r\beta(v_{ce})} + (N - 1) F_{ce}(v_{ce})^{N-2} f_{ce}(v_{ce}) e^{r\beta(v_{ce})}.$$

Substituting this expression for $\zeta'(v_{ce})$ into the prior equation, we obtain

$$(N - 1) F_{ce}(v_{ce})^{N-2} f_{ce}(v_{ce}) = \zeta'(v_{ce}) e^{-rv_{ce}}.$$

Solving for $\zeta'(v_{ce})$, we have

$$\zeta'(v_{ce}) = e^{rv_{ce}} (N - 1) F_{ce}(v_{ce})^{N-2} f_{ce}(v_{ce}) = e^{rv_{ce}} \frac{d}{dv_{ce}} F_{ce}(v_{ce})^{N-1},$$

and noting that $\zeta(0) = 0$, alternative expression for $\zeta(\theta)$ is then given by

$$\zeta(v_{ce}) = \int_{\rho=0}^{v_{ce}} e^{r\rho} dF_{ce}(\rho)^{N-1}.$$

We can now derive the desired expression for the bidding function by substituting this into the expression defining $\zeta(v_{ce})$. This gives us

$$\begin{aligned}\beta(v_{ce}) &= \frac{1}{r} \ln \left[\frac{\zeta(v_{ce})}{F_{ce}(v_{ce})^{N-1}} \right] \\ &= \frac{1}{r} \ln \left[\frac{1}{F_{ce}(v_{ce})^{N-1}} \int_{\rho=0}^{v_{ce}} e^{r\rho} dF_{ce}(\rho)^{N-1} \right] \\ &= \frac{(1-N)}{r} \ln(F_{ce}(v_{ce})) + \frac{1}{r} \ln \left[\int_{\rho=0}^{v_{ce}} e^{r\rho} dF_{ce}(\rho)^{N-1} \right].\end{aligned}$$

We now integrate by parts, letting $dv = dF_{ce}(\rho)^{N-1} \implies v = F_{ce}(\rho)^{N-1}$ and $u = e^{r\rho} \implies du = re^{r\rho}d\rho$. Making the appropriate substitutions gives us

$$\begin{aligned}\beta(v_{ce}) &= \frac{(1-N)}{r} \ln(F_{ce}(v_{ce})) + \frac{1}{r} \ln \left[e^{r\rho} F_{ce}(\rho)^{N-1} \Big|_{\rho=0}^{v_{ce}} - \int_0^{v_{ce}} F_{ce}(\rho)^{N-1} e^{r\rho} r d\rho \right] \\ &= \frac{(1-N)}{r} \ln(F_{ce}(v_{ce})) + \frac{1}{r} \ln \left[e^{rv_{ce}} \left\{ F_{ce}(v_{ce})^{N-1} - \int_0^{v_{ce}} F_{ce}(\rho)^{N-1} e^{-r(v_{ce}-\rho)} r d\rho \right\} \right] \\ &= \frac{(1-N)}{r} \ln(F_{ce}(v_{ce})) + v_{ce} + \frac{1}{r} \ln \left[F_{ce}(v_{ce})^{N-1} - \int_0^{v_{ce}} F_{ce}(\rho)^{N-1} e^{-r(v_{ce}-\rho)} r d\rho \right] \\ &= v_{ce} + \frac{1}{r} \ln \left[1 - \frac{re^{-rv_{ce}}}{F_{ce}(v_{ce})^{N-1}} \int_0^{v_{ce}} F_{ce}(\rho)^{N-1} e^{r\rho} d\rho \right].\end{aligned}$$

Finally, re-introducing the dependence on θ and the subscript as well as specifying a particular broker n , we have our result

$$\beta_S(\theta_n) = v_{ce}(\theta_n) + \frac{1}{r} \ln \left[1 - \frac{re^{-rv_{ce}(\theta_n)}}{F_{ce}(v_{ce}(\theta_n))^{N-1}} \int_0^{v_{ce}(\theta_n)} F_{ce}(\rho)^{N-1} e^{r\rho} d\rho \right].$$

□

While the above discussion has focused on the derivation of the unique pure strategy symmetric Bayesian-Nash equilibrium bidding strategy, it turns out that this is also essentially the unique pure strategy Bayesian-Nash equilibrium. In particular, as the following result establishes, if there is an asymmetric pure strategy Bayesian-Nash equilibrium, the bidding function must be equal to that of the pure strategy symmetric Bayesian-Nash equilibrium except possibly for a set of types of measure zero.

Theorem 5.3. *Let Assumptions 4.1 and 4.2 hold and assume that the distribution of broker types θ_n has a density. In a standard first-price auction, there exists a unique pure strategy Bayesian-Nash equilibrium bidding function up to a set of types of measure zero.*

Proof. It may be shown that our model satisfies the assumptions of Theorem 1 in McAdams (2007), which establishes existence and uniqueness of a pure strategy Bayesian-Nash equilibrium. \square

Uniqueness of the equilibrium bidding function increases confidence in our derived Bayesian-Nash equilibrium as a model of broker behavior.

5.2 Intermediated First-Price Auction

In the intermediated first-price auction, a trusted party solicits from the seller a description of the portfolio and from each broker a function that specifies his bid contingent on the portfolio. In the context of our model, the portfolio's description can be thought of as the parameter q and each broker's bid function β_I^n is a function of q and the number of brokers N . Note that unlike in the standard auction, here N is known to the intermediary and may be reflected in bid values without the requirement that its value be known to brokers, which we do not assume. The intermediary assigns the portfolio to the broker who offers the highest bid $\beta_I^n(q, N)$ and the portfolio is purchased at this price.

The n th broker's choice of bid function β_I^n should depend on his type θ_n . Alternatively, we can view the bid function as a function of three variables, with $\beta_I^n(\theta_n, q, N)$ being the bid the broker assigns given his value of θ_n , the portfolio parameter q , and the number of brokers N . With this perspective, what the broker supplies to the intermediary is a function $\beta_I^n(\theta_n, \cdot, \cdot)$ of the portfolio parameter and broker number. We will view the strategy employed by each broker as such a function of three variables. As a solution concept, we assume as before that brokers behave according to a pure strategy Bayesian-Nash equilibrium. It is easy to extend results from the previous section to show that there exists a unique pure strategy Bayesian-Nash equilibrium, which is symmetric. We therefore drop the superscript n . In particular, for each n ,

$$\beta_I(\theta_n, q, N) \in \operatorname{argmax}_{b \in \mathfrak{R}} E \left[u \left((v_n - b) \mathbf{1}(b \geq \max_{m \neq n} \beta_I(\theta_m, q, N)) \right) \mid \theta_n, q, N \right].$$

The following result serves a purpose analogous to Theorem 5.2 but in the context of an intermediated auction.

Theorem 5.4. (Bayes-Nash Equilibrium Bids in an Intermediated First-Price Auction) *Let Assumptions 4.1 and 4.2 hold and assume that the distribution of broker types θ_n has a density. In an intermediated first-price auction, the unique pure strategy Bayesian-Nash equilibrium bidding function up to a set of types of measure zero is given by*

$$\beta_I(\theta_n, q, N) = v(\theta_n) + \frac{1}{r} \ln \left[1 - \frac{r e^{-rv(\theta_n)}}{F_v(v(\theta_n))^{N-1}} \int_0^{v(\theta_n)} F_v(\rho)^{N-1} e^{r\rho} d\rho \right], \quad (3)$$

where $v(\theta_n) = v^* - \theta_n q$ is a function mapping broker types to broker valuations and F_v is the cumulative distribution of these valuations.

Proof. The proof is analogous to that of Theorem 5.2 but now making use of the full information available to the intermediary when calculating brokers' bids. Uniqueness again follows from Maskin and Riley (1984) and McAdams (2007). \square

6 Benefits of an Intermediary

Based on the derived Bayesian-Nash equilibria bidding strategies for the standard and intermediated first-price auctions derived in Section 5, several numerical studies were performed to assess how these two mechanisms compare regarding the seller's transaction costs. This was done via Monte Carlo averaging over specific values for q and θ_n drawn from their respective distributions. In each case, the constants characterizing a given study are N and r . In the following subsections we first discuss the motivations behind our choices for the specific parameter values used in the numerical studies and then present our results.

6.1 Representative Instances

We will discuss results of computations carried out to assess the benefits afforded through use of an intermediary. These computations involve model instances with specific parameter settings. In particular, we make Assumptions 4.1 and 4.2 and consider instances in which $v^* = 5 \times 10^8$, $q \sim \text{unif}([0, 10^7])$, and $\theta_n \sim \text{unif}([0, 1])$. We consider ranges for the number of brokers $N \in \{2, \dots, 12\}$ and the risk aversion parameter $r \in [0, 4 \times 10^{-7}]$. Let us motivate these choices.

The values of v^* and the ranges for q , and θ_n are chosen to reflect a realistic portfolio value of half a billion dollars and a realistic range from zero to ten million dollars of potential broker transaction costs. Blind portfolio auctions typically engage two to four dominant brokers. The range we consider for N includes these possibilities and extends beyond them.

The choice of risk aversion parameter deserves further discussion. The case of risk-neutrality ($r = 0$) serves as one extreme worth understanding for insight though it is unrealistic. Our intention is for realistic risk aversion parameter values to be captured within the range $[0, 4 \times 10^{-7}]$. To understand why this intention should be served, consider a situation where a broker is uncertain about what a portfolio is worth to him and assumes a normal distribution with a standard deviation of one million dollars around his expectation. What premium would have to be subtracted from his expectation to arrive at a price where he is indifferent about acquiring the portfolio? Different brokers would have different opinions on this matter, but a representative figure might be one hundred thousand dollars, which is ten percent of the standard deviation. This guess implies a risk aversion of $r = 2 \times 10^{-7}$. In particular, letting the broker's profit, which is the difference between the amount he pays for the portfolio and the amount he later discovers it is worth to him, be denoted by $x \sim N(10^5, (10^6)^2)$, this value of r solves

$$v_{\text{CE}}(\theta_n) = -\frac{1}{r} \ln(-E[-e^{-rx}]) = 10^5 - \frac{r}{2}(10^6)^2 = 0.$$

6.2 Benefits to the Seller

Prior to participating in the auction, the seller holds a portfolio with a spot value of v^* , i.e. the portfolio value at the time at which the trade price is being fixed, such as the market closing price, etc. Note that this nominal value is assumed known by each of the N brokers, i.e. common market knowledge. Through participation in the auction, the seller is able to liquidate this portfolio by selling it to the winning broker, whose bid we will denote by b_W . When liquidating the portfolio, the seller's goal is to do so while minimizing the transaction costs incurred in the process. Letting T denote the transaction costs of the seller, we have that $T = v^* - b_W$, representing the difference between the "book value" of the portfolio and what the seller receives for it.

As our model assumes a risk-neutral seller, the metric of concern from the seller’s point of view is the expected transaction cost $E[T]$. Computing the brokers’ Nash equilibrium bidding strategies using the expressions derived in Section 5, we are able to numerically compute $E[T]$ and study how the standard and intermediated first-price auctions compare as we vary two parameters, the brokers’ common risk aversion parameter r , as well as the number of brokers N .

The importance of studying the change with respect to r is immediate in that this gauges how relevant the informational asymmetry between the seller and brokers is as seen from the brokers’ point of view and how this in turn affects the seller. The number of brokers N from which the seller receives bids is important in that it represents the only degree of freedom, or operating parameter, left to the discretion of the seller once he has chosen either the standard or intermediated mechanism. As mentioned in Section 1, typically sellers collect bids from a small handful of brokers, often on the order of $2 \sim 4$, in order to avoid information leakage. By introducing a trusted intermediary, sellers are safely able to solicit bids from an arbitrarily large number of brokers without risk. Given this possibility, it is important to consider the potential benefits of higher values of N made possible by an intermediary and compare this to the standard case.

Consider Figures 2 and 3, which compare the average seller transaction cost in standard and intermediated first price auctions as r and N are varied, respectively. There are several points that summarize our findings. In what follows, we will use $\bar{T}(A)$ to denote the average seller transaction cost when the auction is of type $A \in \{S, I\}$, denoting a standard or intermediated auction.

We see that across all risk aversion levels $r > 0$ that $\bar{T}(I) < \bar{T}(S)$ and that the gap between the two increases with r . There are in essence two pressures driving broker behavior as risk aversion grows. On the one hand, increased risk aversion leads to the possibility for additional expected utility gains from bidding higher, resulting from marginal probability of winning gains that outweigh marginal utility losses when the bid is increased. This tends to push the expected seller revenue up and hence improves the seller’s transaction cost. However, as r increases, v_{ce} for each broker decreases as well, resulting in lower bids, lower expected seller revenues, and hence higher transaction costs. Which force dominates will

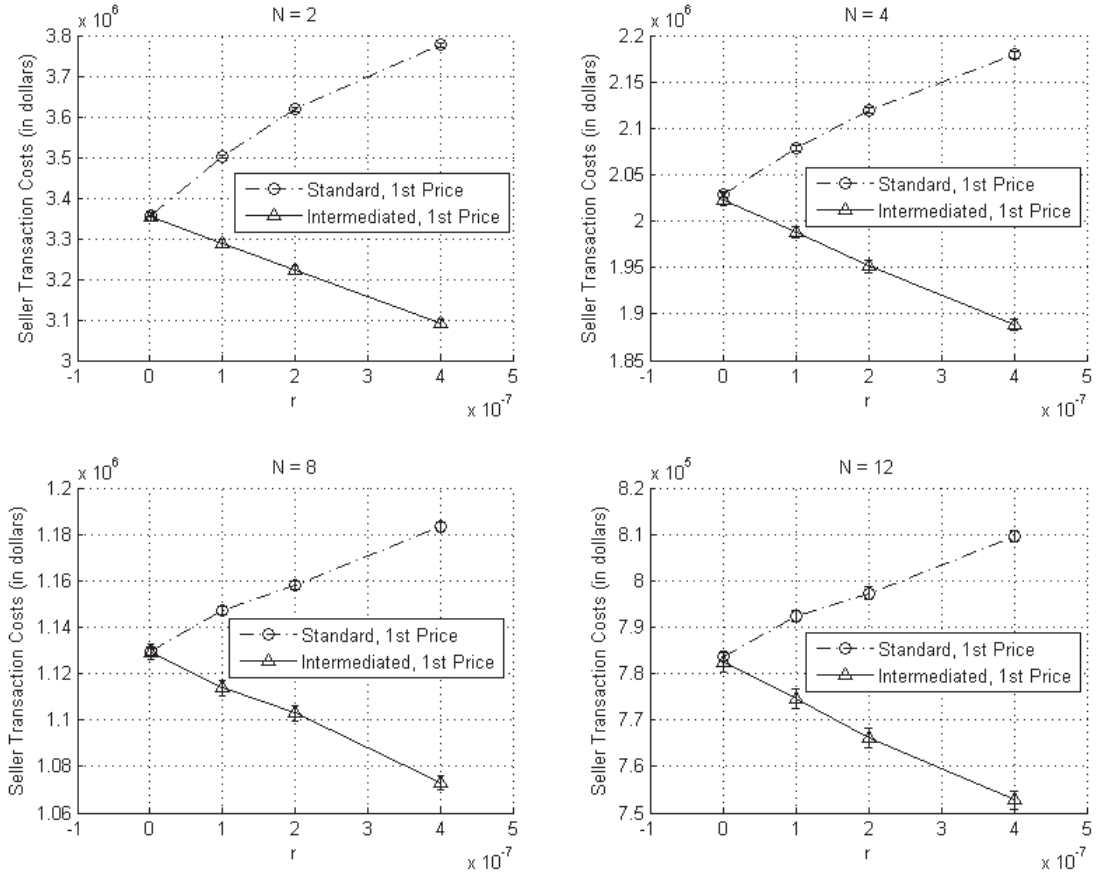


Figure 2: Average Seller Transaction Cost vs. Risk-Aversion (95% Confidence Intervals)

depend on the degree of uncertainty in the portfolio being auctioned and this will determine how increasing r affects the seller's transaction costs. In the case of the standard auction, the portfolio uncertainty dominates and hence increasing r increases $\bar{T}(S)$.

Interestingly, the other pressure dominates in the intermediated case and we see that $\bar{T}(I)$ decreases as r increases. To explore this, consider an intermediated auction under Assumption 4.2. We will refer to one utility function u_2 as being more risk averse than a second u_1 if it can be represented as a concave transformation of the second, i.e. $u_2 = \zeta(u_1)$, where ζ is an increasing concave function. Let $r_A(x, u)$ denote the Arrow-Pratt coefficient for absolute risk aversion of a utility function u at x and let $c(F, u)$ denote the certainty equivalent of a gamble F under utility function u . Then it may be shown (Mas-Colell et al. (1995)) that this relationship between utility functions is equivalent to each of the individual

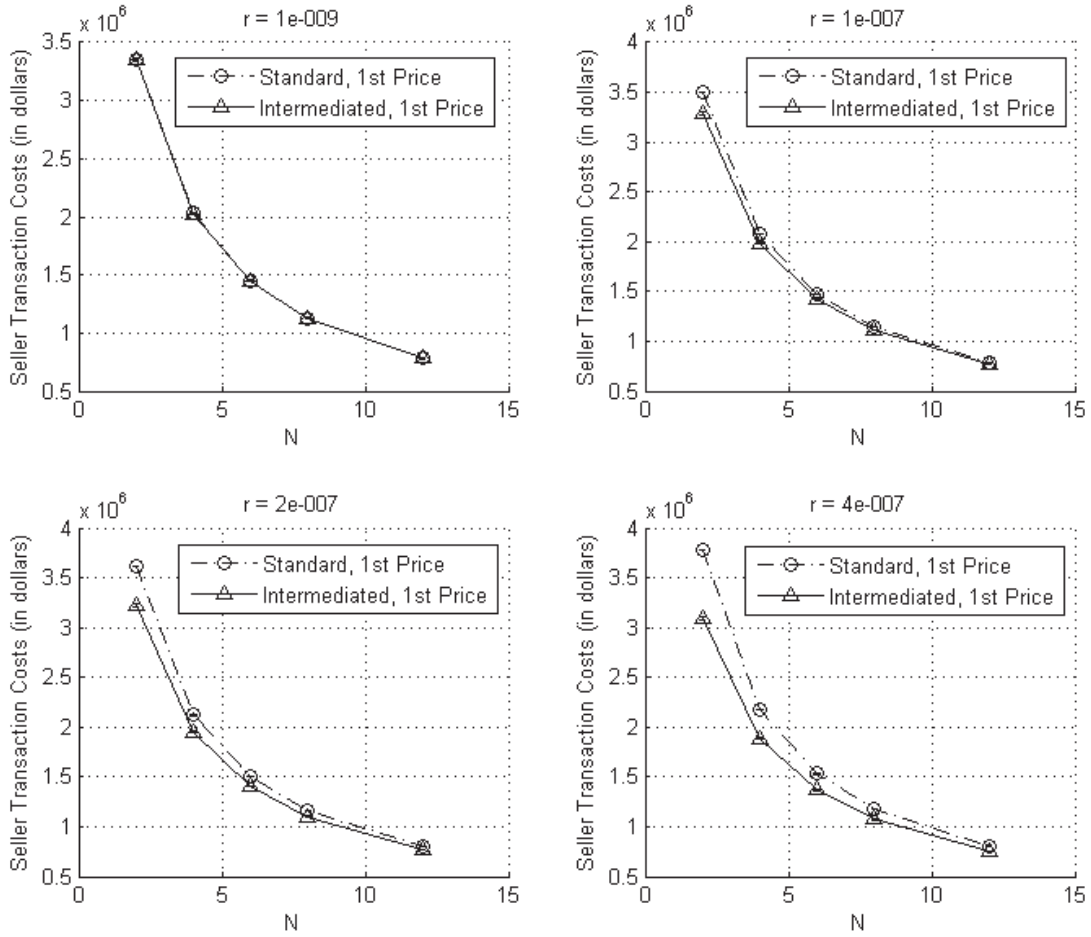


Figure 3: *Average Seller Transaction Cost vs. Number of Brokers (95% Confidence Intervals)*

statements that $r_A(x, u_2) \geq r_A(x, u_1) \forall x$ and $c(F, u_2) \leq c(F, u_1)$ for any F . In addition, we define U_0 to be the set of differentiable monotonically increasing utility functions such that $u(0) = 0$. The following theorem establishes in a fairly general setting that seller transaction costs decrease as broker risk aversion increases, as observed in our numerical study. Although we focus on utility functions in U_0 in this result, note that it applies to any family of utility functions that may be transformed into U_0 via a constant offset as this is irrelevant to the study of relative mechanism dynamics. This thus applies to the utility form adopted in Assumption 4.1.

Theorem 6.1. (Transaction Costs Decrease as Risk Aversion Increases in an Intermediated

Auction) Let Assumption 4.2 hold. For all $u_1, u_2 \in U_0$ such that u_2 is more risk averse than u_1 , $\bar{T}(I)$ is strictly less if brokers realize utility u_2 than if brokers realize utility u_1 .

Proof. Note that an intermediated auction is conceptually equivalent to a first-price auction where q is common knowledge among all brokers and each broker bids optimally given that information. We proceed under this framework. Assume that in equilibrium there exists a monotonically increasing and differentiable bidding function β mapping values to bids with the property that $\beta(0) = 0$. Focusing on a given broker with value v , if all other brokers follow the equilibrium strategy β then the broker solves for their optimal bid value according to

$$\max_z H(z)u(v - \beta(z)),$$

where z functions as the broker's decision variable when determining his bid and H is the distribution of the highest value of the other $N - 1$ players, also assumed differentiable. Taking the first derivative with respect to z gives us

$$h(z)u(v - \beta(z)) - H(z)u'(v - \beta(z))\beta'(z) = 0,$$

or in equilibrium when brokers bid according to their true values,

$$\beta'(v) = \frac{h(v)}{H(v)} \frac{u(v - \beta(v))}{u'(v - \beta(v))}. \quad (4)$$

Now consider two different risk-averse utility functions $u_1, u_2 \in U_0$ such that u_2 represents a greater degree of risk aversion, i.e. we may relate the two via $u_2 = \zeta(u_1)$, where ζ is an increasing concave function. Using (4), we have

$$\begin{aligned} \beta'_2(v) &= \frac{h(v)}{H(v)} \frac{u_2(v - \beta_2(v))}{u'_2(v - \beta_2(v))} \\ &= \frac{h(v)}{H(v)} \frac{\zeta[u_1(v - \beta_2(v))]}{\{\zeta[u_1(v - \beta_2(v))]\}' } \\ &= \frac{h(v)}{H(v)} \frac{\zeta[u_1(v - \beta_2(v))]}{\zeta'[u_1(v - \beta_2(v))]u'_1(v - \beta_2(v))}. \end{aligned} \quad (5)$$

From our assumption that $u_1, u_2 \in U_0$, implying $u_1(0) = u_2(0) = 0$, it follows necessarily that $\zeta(0) = 0$. Using this fact along with the concavity of ζ , it may be easily shown via a first-order Taylor expansion that $\zeta(x)/\zeta'(x) > x$ for all x in its domain. Applying this

inequality to (5) gives us

$$\begin{aligned}\beta'_2(v) &= \frac{h(v)}{H(v)} \frac{\zeta[u_1(v - \beta_2(v))]}{\zeta'[u_1(v - \beta_2(v))u'_1(v - \beta_2(v))]} \\ &> \frac{h(v)}{H(v)} \frac{u_1(v - \beta_2(v))}{u'_1(v - \beta_2(v))}.\end{aligned}\tag{6}$$

Now note that (4) may be also applied to u_1 to give us

$$\beta'_1(v) = \frac{h(v)}{H(v)} \frac{u_1(v - \beta_1(v))}{u'_1(v - \beta_1(v))}.\tag{7}$$

Consider now the scenario in which $\beta_1(v) > \beta_2(v)$. This leads to the relationship that $v - \beta_2(v) > v - \beta_1(v)$. Since u_1 is defined to be concave, $u_1(x) \uparrow$ and $u'_1(x) \downarrow$ as $x \uparrow$. Hence, we may conclude that for $\beta_1(v) > \beta_2(v)$ that

$$\beta'_2(v) > \frac{h(v)}{H(v)} \frac{u_1(v - \beta_2(v))}{u'_1(v - \beta_2(v))} > \frac{h(v)}{H(v)} \frac{u_1(v - \beta_1(v))}{u'_1(v - \beta_1(v))} = \beta'_1(v).$$

Thus, we have shown that $\beta_1(v) > \beta_2(v) \Rightarrow \beta'_2(v) > \beta'_1(v)$. Combining this with the knowledge that $\beta_1(0) = \beta_2(0) = 0$, we have that $\beta_2(v) > \beta_1(v) \forall v > 0$. Since for any given value v , brokers with strictly greater risk-aversion will compute strictly larger equilibrium bids, it follows that the corresponding seller revenue will be strictly greater. Finally, because the seller's transaction costs decrease with increasing revenue, it follows that $\bar{T}(I)$ strictly decreases. \square

The intuition behind Theorem 6.1 may be seen in Figure 4. In the top plot we see two utility functions, one more risk-averse (dashed curve) than the other (solid curve). Without any change in broker behavior, we may shift and scale the dashed curve, as seen in the bottom plot, to start from the origin and cross the solid curve at the point corresponding to the solid curve's utility when a broker in that scenario wins the auction by bidding the optimal amount $\beta(v)$ when their valuation is v . When normalized this way, the slope of the less risk-averse utility function is always greater than that of the more risk-averse utility function at that point. As a result, the optimal bid for the more risk-averse case must be greater than $\beta(v)$ as the marginal decrease in expected utility as the bid increases is less, i.e. the optimal point where the marginal decrease in expected utility balances the marginal gain in winning probability will occur for a larger optimal bid.

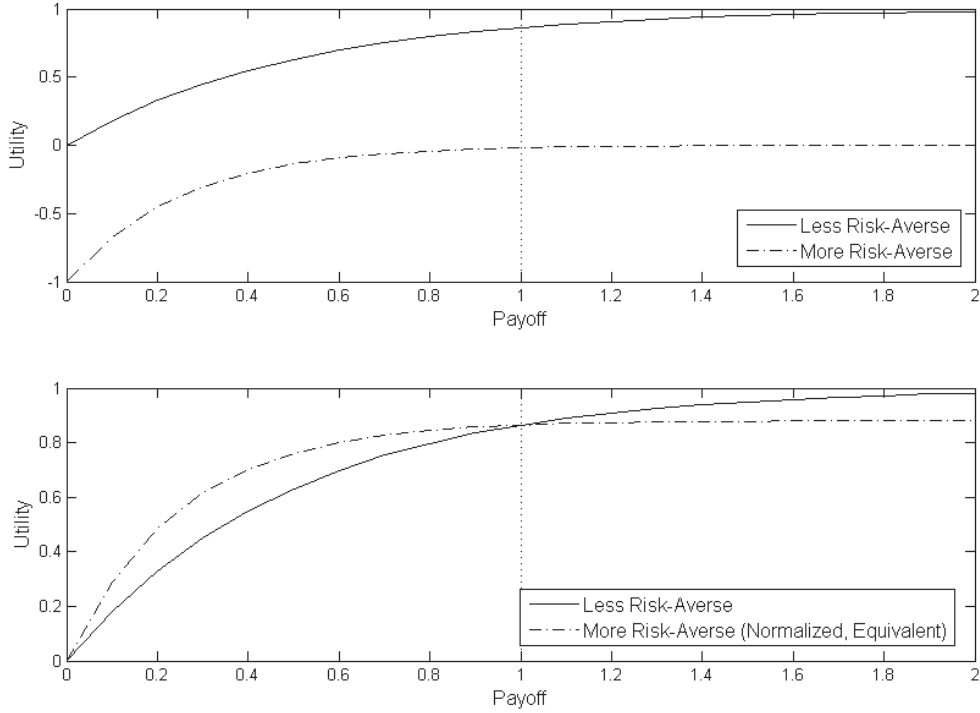


Figure 4: *Two utility functions, the more risk-averse one normalized in the second figure such that $u_2(0) = 0$ and $u_1(v - \beta(v)) = u_2(v - \beta(v))$. Here $v - \beta(v) = 1$ corresponds to the optimal bid $\beta(v)$ for the less risk-averse utility function.*

Interpreting these results in terms of representative parameters discussed in Section 6.1, $r = 2 \times 10^{-7}$ and the inclusion of an intermediary provides the seller with a transaction cost savings of over 10% relative to the standard auction for the case of $N = 2$. This represents a significant improvement for institutional investors.

With respect to changes in the number of brokers N , it is observed that for both the standard and intermediated cases and for all $r > 0$ that the seller's transaction costs are convex in N and asymptotically approach a common fixed lower bound as N increases. This result is driven by two factors. First, as N increases, $\min_n \theta_n \rightarrow 0$, and hence in both the standard and intermediated cases the value of the winning broker approaches v^* . Second, as N increases the degree to which a broker in a first-price auction will bid below his valuation decreases in order to remain competitive. In the limit, brokers in both auctions bid asymptotically closer to v^* as N increases. It is the ability of the intermediated case to

prevent information leakage that allows sellers to increase N to approach this limit arbitrarily closely, minimizing their transaction costs.

7 Intermediary Implementation

An important issue is how to implement the notion of a trusted intermediary. In order for the auction mechanism to function properly, participants need to be guaranteed that the information revealed to the intermediary (i.e. bidding functions $\beta_I(\theta_n, \cdot)$ from brokers and portfolio characteristic q from the seller) will be kept completely confidential and used appropriately. If that is not the case, then the proposed system and its merits will decay due to two factors. First, brokers will become increasingly hesitant to participate in the market if they feel that their information may be compromised, denying both the seller and brokers the opportunity to potentially gain from a principal trade. The second result of a lack of trust in the intermediary is that risk-averse brokers will factor the potential for information leakage into their bids, which will be subsequently worse for the seller.

One possible form for an intermediary would be a government sanctioned organization, perhaps itself a division of the SEC, that has sufficient transparency and auditory protocols in place to ensure information security. This could also be a role potentially made available to non-governmental third party companies, and competition amongst them could lead to the evolution of a protocol for handling these transactions that best satisfies participants' needs. Important would be the ability for participants to verify the results of auctions when necessary, and various details regarding issues like participant anonymity, etc. would need to be carefully considered. Note that the finance industry has already seen the advent of third party companies that facilitate portfolio trades, though these companies generally get involved in sizable transition trades where a seller is making a large shift in assets or perhaps a pension fund has changed management.

A second manner in which the intermediary may be implemented is via an equivalent cryptographic protocol. In the past decade, researchers in computer science have begun looking at ways in which concepts from cryptography, namely zero-knowledge proofs and related protocols, may be applied to various settings in electronic commerce, auctions, and

security exchanges. In Thorpe and Parkes (2009), for instance, a methodology is presented by which a seller's order may be completely filled, making use of a secure and verifiably accurate way to report to potential brokers the risk parameters of a portfolio formed by combining their current positions and those of the seller, with neither party knowing the positions of the other. In addition, Izmalkov et al. (2008), Micali (2010), and Izmalkov et al. (2011) present theoretical work that develops methodologies by which mechanisms and trusted intermediation may be implemented by the agents themselves in a manner that is verifiable and without anyone's private information being divulged. While currently limited in practicality when applied to arbitrarily complicated mechanisms, this line of research suggests that further work will allow for a distributed cryptographic implementation of an intermediary in the future.

8 Conclusion

This research has studied and demonstrated the reduction in seller transaction costs made possible by incorporating an intermediary between sellers and competing brokers in the standard blind portfolio auction protocol. Using a model designed to focus primarily on the notion of uncertainty faced by auction participants, this work has demonstrated via both simulation and theoretical results the potential benefit to portfolio sellers of the use of an intermediary and the associated removal of uncertainty from the process. For example, it was shown that under reasonable assumptions on brokers' utility functions and investment preferences, it is possible to realize seller transaction cost improvements of over 10%. This degree of savings is very significant to many institutional investors.

To our knowledge, this is the first investigation of how the notion of uncertainty affects the efficiency of blind portfolio auctions. With a daily shares traded percentage of roughly 12% on the NYSE, and similar numbers on other major exchanges, it is important that this prominent form of business transaction be better understood so as to make more efficient use of the considerable capital represented by these trades.

In future work, we intend to enhance our model to include factors allowing us to account for the allocative inefficiency brokers face in blind portfolio auctions in practice. As discussed

in Section 1, a large component of how brokers value a portfolio ex-post that they have won is based on idiosyncratic factors such as their pre-existing positions, proprietary trading strategies, etc. When portfolios are auctioned blindly, it will sometimes be the case that the winning broker is not the one who most valued the portfolio and thus the market allocates inefficiently. The effect is that brokers will reduce their bids to hedge themselves against this risk and hence not only are brokers not extracting as much surplus from the trade as they could, but seller transaction costs are also hurt by lower broker bids. The introduction of a trusted intermediary should thus not only improve the auction's performance from the point of view of sellers, but the overall social welfare of all auction participants as well for a Pareto improvement in moving from the standard to intermediated auction.

References

- R. Almgren and N. Chriss. Bidding principles. *Risk*, June 2003.
- N.P.B. Bollen, R.E. Whaley, and T. Smith. Modeling the bid/ask spread: Measuring the inventory-holding premium. *Journal of Financial Economics*, 72, 2004.
- J. Ferry. Peak time for programme trading. *Risk*, May 2004.
- C. Giannikos and M. Suen. Pricing determinants of blind principal bidding and liquidity provider behavior. In *European Financial Management Meeting*, June 2007a.
- C. Giannikos and M. Suen. Estimating two structural spread models for trading blind principal bids. In *European Financial Management Meeting*, June 2007b.
- C. Giannikos, H. Guirguis, and T.S. Suen. Modeling the blind principal bid basket trading cost. *European Financial Management*, December 2009.
- S. Izmalkov, M. Lepinski, and S. Micali. Verifiably secure devices (and correlated equilibrium). In *Theory of Cryptography Conference*, February 2008.
- S. Izmalkov, M. Lepinski, and S. Micali. Perfect implementation. *Games and Economic Behavior*, 71(1):121–140, January 2011.

- K. Kavajecz and D. Keim. Packaging liquidity: Blind auctions and transaction cost efficiencies. *Journal of Financial and Quantitative Analysis*, 40(3), 2005.
- D. Keim and A. Madhavan. The upstairs market for large block transactions: Analysis and measurement of price effects. *The Review of Financial Studies*, 9(1), 1996.
- A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- E. Maskin and J. Riley. Optimal auctions with risk averse buyers. *Econometrica*, 52(6), November 1984.
- D. McAdams. Uniqueness in symmetric first-price auctions with affiliation. *Journal of Economic Theory*, September 2007.
- S. Micali. Perfect concrete implementation of arbitrary mechanisms: A quick summary of joint work with Sergei Izmalkov and Matt Lepinski. In *Behavioral and Quantitative Game Theory: Conference on Future Directions*, pages 1–5. ACM, 2010.
- P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.
- G. Sofianos. Dark pools and algorithmic trading. In *Algorithmic Trading Handbook, 2nd Edition*. The Trade, 2007.
- H.R. Stoll. The supply of dealer services in securities markets. *Journal of Finance*, 33(4), 1978a.
- H.R. Stoll. The supply of dealer services: An empirical study of NASDAQ stocks. *Journal of Finance*, 33(4), 1978b.
- T.S. Suen. *Essays on Modeling of Blind Principal Bid Basket Trading Cost*. PhD thesis, The City University of New York, 2007.
- C. Thorpe and D. Parkes. Cryptographic combinatorial securities exchanges. In *Financial Cryptography and Data Security (FC '09)*, 2009.
- Traders Magazine. Principal blind bidding in portfolio trading. (*Special to*) *Traders Magazine*, August 1998.