Trading in Networks: A Normal Form Game Experiment*

Douglas Gale[†]
New York University

Shachar Kariv[‡]
UC Berkeley

March 22, 2007

Abstract

This paper reports an experimental study of trading networks, in which exchange is intermediated by traders who form a chain of links between the initial owner of the assets and ultimate owner of the assets. Traders choose bid and ask prices and trades are executed by the computer once subjects have submitted their strategies. Networks are incomplete in the sense that each trader can only exchange assets with a limited number of other traders. The greater the incompleteness of the network, the more intermediation is required to transfer the assets between initial and final owners. The uncertainty of trade in networks constitutes a potentially important market imperfection. As a result, the inferences subjects must draw in order to make optimal decisions are quite subtle. Nevertheless, we find that the competitive prices can account for the pricing behavior observed in the laboratory in variety of networks and trading protocols. Furthermore, significant differences can be identified in the pricing behavior of subjects in different networks, and different trading protocols lead to different dynamics.

Journal of Economic Literature Classification Numbers: C91, C92, G10, G19.

*This research was supported by the Center for Experimental Social Sciences (C.E.S.S.) at New York University. We are grateful to Peter Bossaerts, Gary Charness, Syngjoo Choi, Tom Palfrey, Charles Plott, and Bill Zame for helpful discussions. Kariv is grateful to the hospitality of the Institute for Advances Studies School of Social Science.

[†]Department of Economics, New York University, 269 Mercer Street, #507, New York, NY, 10003 (E-mail: douglas.gale@nyu.edu, URL: http://www.econ.nyu.edu/user/galed).

[‡]Department of Economics, University of California, Berkeley, Evans Hall # 3880, Berkeley, CA 94720 (E-mail: kariv@berkeley.edu, URL: http://socrates.berkeley.edu/~kariv/).

1 Introduction

Markets play a central role in economic theory and practice, but economic models often give short shrift to the important institutional details of markets. For example, costly intermediation ("frictions") plays a crucial role in many financial markets and yet the standard model is a centralized auction market where everyone trades simultaneously at a single location. While stock exchanges and other financial markets are usually assumed to be good approximations to the theoretical ideal of frictionless markets, recent theoretical research has begun to recognize that financial markets contain interesting and important frictions. Financial networks, which are crucial for the allocation of resources in society, are a natural example to study, but results are applicable to any model of exchange which shares the same basic network structure.

Gale and Kariv (2007) contribute to a more realistic theory by developing a model of financial networks. The network is represented by a connected graph in which the nodes represent agents (traders) and the edges represent the possibility of trade between the agents linked by the edge. When the network is incomplete, it may be necessary for an asset to follow a circuitous route from the initial seller to the ultimate buyer. The more incomplete the network, the more intermediate trades are required to achieve an efficient allocation. Gale and Kariv (2007) show that, in the limit as the economy becomes frictionless, the market outcome is efficient. Introducing frictions provides an important source of market imperfection which can lead to a market breakdown.

In this paper, we exploit the methods of experimental economics to explore the properties of a simple trading network. Empirical research can tap either real-world data from large-scale markets or small-scale laboratory data. The strengths of data from the real world are its relevance and availability. Its main weakness is that in real-world settings we observe behavior but not preferences, technologies, or private information. In the laboratory, by contrast, we can control subjects' preferences, technology and private information. Consequently, laboratory data are especially useful for testing the efficiency of different markets, and for comparing market structures and market institutions. Thus, the clarity that is achieved by putting a market under the microscope is well worth the effort and the necessary simplification.

The theory involves a number of different elements, each of which raises questions about subjects will behave in the laboratory. Will subjects bargain as the theory suggests? Will subjects behave rationally or will it be necessary to allow for bounded rationality? Will intermediaries' uncertainty about the possibility of reselling the asset lead to a coordination failure? In this paper, we study a market in which the number of traders is small, the networks are simple, and the trading mechanism is closer to the well known auction paradigm than to the bargaining paradigm. The design is stable and easy to understand and provides us with an insight into how experimental networks behave. Using this design, we test how useful the theory is in interpreting the observed behavior, and study the efficiency of pricing and trade using a variety of network architectures.

A parametric example may clarify the experimental design. Suppose there are nine subjects arranged in the rectangular array with three rows and three columns illustrated in Figure 1. In addition to the human traders, there is a computer-generated seller (CGS) and a computer-generated buyer (CGB). Each node represents a trader and the edges between the nodes indicate trading possibilities. The network architecture in Figure 1 indicates that trades are restricted to adjacent rows but, subject to these constraints, any pattern of trading links is allowed. That is, each member of the top row can trade with the CGS and with every member of the middle row; each member of the middle row can trade with every member of the top and bottom rows; and each member of the bottom row can trade with every member of the middle row and with the CGB.

[Figure 1 here]

The CGS is endowed with a single unit of an indivisible asset. The nine traders are endowed with 100 tokens each. Buyers use these tokens to pay for the asset and sellers receive these tokens in exchange for the asset. The CGB is also assumed to have an endowment of 100 tokens. The asset has no value to the CGS or to the nine traders. The CGB values the asset at 100 tokens. So the surplus (gains from trade) generated by transferring the asset from the CGS to the CGB is equal to 100 tokens. Each trader simultaneously chooses a bid (the price at which he is willing to buy the asset) and an ask (the price at which he is willing to sell the asset). The bids and asks must lie between 0 and 100 tokens. The ask of the CGS is fixed at 0 and the bid of the CGB is fixed at 100.

Once the bids and asks have been determined, trades are executed as follows. Beginning at the top of the network, the CGS and the top row exchange the asset. The asset goes to the trader with the highest bid. If two or more traders choose the highest bid, the asset is allocated randomly between them (with equal probabilities). The top-row seller (the trader who bought the asset from the CGS) sells the asset to the middle-row trader with the highest bid that is at least as high as the seller's ask. Again, ties are broken randomly. If every bid is less than the seller's ask, no trade takes place and the game ends with the seller holding the asset. Exchange between the middle and bottom rows is executed similarly. Finally, if the asset reaches the bottom row, the asset will be transferred to the CGB because the CGB's bid of 100 is at least as great as the seller's ask. When the asset is traded, the price paid for the asset is the average of the bid and the ask. The corresponding amount of tokens is transferred from the buyer to the seller.

This example gives a good sense of the advantages of the experimental design. First, it defines a normal form game. Secondly, because the game is played in normal form, it can be played repeatedly in a relatively short amount of time, generating a large data set. Thirdly, the platform is sufficiently flexible to allow us to study a variety of network architectures, transaction pricing rules, and payoff functions. The baseline treatment (**B**) uses the 2×3 and 3×3. When the asset is traded, the price paid for the asset is the average of the bid and the ask. The corresponding amount of tokens is transferred from the buyer to the seller. At the end of an experimental session, one trading period is chosen at random to determine the subjects' payoffs. A subject's earnings in this period equal his initial endowment of 100 tokens plus his trading profit (positive or negative). The subject's payment is equal to the greater of his earnings minus 90 tokens and 10 tokens. Thus, the payoff function does not deduct subjects' full trading losses from their earnings. The experimental design section discusses the variety of other treatments that are available.

Our results can be summarized under three headings:

- *Convergence*. Since the underlying trading game is essentially a Bertrand pricing model, there is a unique Nash equilibrium in which all bid and ask prices are equal to 100 and trading profits are zero, except in the top row, where the CGS is restricted to ask 0 and profits are 50. In the auction treatment (**A**), which uses the 1×3 network and thus corresponds to a simple auction, convergence to the equilibrium price occurs in the first few periods and the bids remain at that level throughout the game, apart from occasional experimental deviations. In the 2×3 and 3×3 networks used in the baseline treatment, convergence is slower but prices very rapidly reach the neighborhood of the competitive price. The slower convergence is due to the greater amount of intermediation required for the asset to reach the CGB.

- *Efficiency.* Strategic uncertainty (about what other subjects will do) inevitably requires a period of learning and during this period trades may not be completed. Further, even later in the game trade may break down if subjects make mistakes about the prices that are likely to be bid or asked by their opponents. On the whole, trade is approximately efficient in the sense that, it tends to be lower in the early trading periods and rises as subjects become more confident about the behavior of other agents and as the prices bid and asked convergence to competitive equilibrium prices. In the 1×3 network, there is no possibility of incomplete trades, since there is only one row. In the 2×3 and and 3×3 networks, efficiency is very high. Given the incompleteness of these networks, which requires intermediate trades, and the strategic form of the game, which does not allow for recontracting, the subjects' ability to coordinate on an efficient outcome is quite striking.

- *Sensitivity.* We study a number of variants to test the sensitivity of the results to the amount of competition (reducing the number of columns in the network), the pricing rule (setting the transaction price equal the buyer's bid), the payoff function (deducting more trading losses from subjects' earnings), and the symmetry of the network architecture (introducing asymmetry in the availability of counterparties to trade with). Although convergence and efficiency are somewhat affected by each of these changes, the equilibrium properties continue to have predictive power. Among other things, we note that (*i*) less competition may lead to slower convergence and lower efficiency, and (*ii*) trading losses and bid-pricing rule reduce competition by making bidders less aggressive and thus lower efficiency.

The rest of the paper is organized as follows. A discussion of the related literature is provided in Section 2. We describe the theoretical model and the experimental design in Section 3. The results are contained in Section 4. Some concluding remarks and important topics for further research are contained in Section 5. Sample experimental instructions are reproduced in Section 6.

2 Related Literature

The paper contributes to the enormous body of work on experimental markets. Following the seminal papers of Forsythe, Palfrey and Plott (1982, 1984), and Plott and Sunder (1982, 1988) numerous experimental papers analyzed many aspects of asset markets.[1] In contrast to the existing literature, our project will develop the first systematic experimental test of the role of network structure in determining the efficiency of markets since the incompleteness of the network provides an important form of market imperfection.

Although network experiments in economics are recent, there is now a large experimental literature on the economics of networks.[2] To the best of our knowledge, all of the previous contributions have quite different focuses than ours. The most closely related paper is by Charness, Corominas-Bosch and Fréchette (2005), who investigate how the network structure affects the outcomes and dynamics of ultimatum bargaining. Following the model of Corominas-Bosch (2004), they decompose a network of buyers and sellers into two simple sub-graphs and test whether it matters how a single edge is added between these two groups of traders.

We provide a couple of fundamental innovations over previous work. Most importantly, previous experimental studies on networks have been restricted to very simple and relatively extreme network architectures. Our primary methodological contribution is an experimental platform that provides a computerized graphical representation of networks and allows subjects to make large numbers of decisions in a wide range of situations. This enables us to systematically collect more and richer data about networks than has been possible in the past. The applications of this platform have not been exhausted by the present paper. The experimental set up can be adapted to the analysis of a larger class of networks and there are many important questions that remain to be explored.

3 Theory and Design

In this section, we describe the theory on which the experimental design is based and the design itself.

[1] Sunder (1995) provides a comprehensive discussion of the experimental work on asset markets.

[2] Kosfeld (2004) surveys the experimental work in economics, and Goyal (2004) and Jackson (2004) provide recent surveys of the theoretical work.

3.1 The trading game

The trading game consists of a finite number of players, indexed by $k = 1, ..., K$, arranged in a rectangular network consisting of m rows and n columns. An example of a 3×3 network is illustrated in Figure 1 above. A single player is located at each node and the edges connecting the nodes indicate that the corresponding players can trade with each other. In addition to the human players, there are two computer-generated players, called the computer-generated seller (CGS) and the computer-generated buyer (CGB). The CGS has one unit of an indivisible asset which he is willing to sell for zero and the CGS is willing to buy the asset for $v > 0$ units. Most of the networks we consider are symmetric and satisfy the following properties. Only the players in row $i = 1$ can purchase the asset from the CGS. The players in row $i > 1$ can buy the asset from any of the players in row $i - 1$ and the players in row $i < m$ can sell the asset to any of the players in row $i + 1$. Only players in the last row can sell the asset to the CGB. A strategy for each player k consists of the announcement of an asking price (a_k) at which he would be willing to sell one unit of the asset and a bid price (b_k) at which he would be willing to purchase one unit of the asset. The bid and ask prices are restricted to the interval $[0, v]$ so the strategy set for player k is simply $S_k = [0, v] \times [0, v]$ and the set of strategy profiles is $S = \times_{k=1}^{K} S_k$.

Trades are executed as follows. The asset is transferred from the CGS to the player row 1 who has the highest bid. If there is more than one player with the highest bid, the asset is allocated randomly among the winning bidders. The player who receives the asset pays the CGS an amount equal to αb_k where b_k is the winning bid and $0 \leq \alpha \leq 1$ is a constant. In each row $1 < i < m$, trade is only possible if the asset is held by one of the players in row $i - 1$, whom we call the seller. Trade takes place if at least one bid in row i is greater or equal to the seller's asking price. The asset is transferred to the highest bidder. If more than one player has the highest bid, the asset is allocated randomly among the winning bidders. The player who receives the asset transfers an amount equal to $\alpha b_k + (1 - \alpha) a_{k'}$, where b_k is the winning bid and $a_{k'}$ is the seller's asking price. If no bid is at least as high as the seller's asking price, no trade takes place and the asset remains with the seller. If a player in row m receives the asset, he sells it to the CGB for a price equal to $\alpha v + (1 - \alpha) a_k$, where a_k is the seller's asking price.

A player's payoff is equal to his trading profit, that is, the amount he receives from selling the asset minus the amount he pays for it. A player who does not manage to buy the asset receives a payoff of zero. A player who buys the asset for a positive price and fails to sell it receives a negative payoff. Denote a typical player's strategy by $\sigma_k = (a_k, b_k)$ and a strategy profile by $\sigma = \{\sigma_k\} = (a, b)$, where $a = \{a_k\}$ and $b = \{b_k\}$. Denote player k's payoff by $\pi_k(\sigma) = \pi_k(a, b)$. A Nash equilibrium is a strategy profile σ^* such that for any player k,

$$\pi_k(\sigma^*) \geq \pi_k\left(\sigma_k, \sigma^*_{-k}\right)$$

for any $\sigma_k \in S_k$. It is not hard to see that the usual Bertrand competition result holds. Suppose that there are at least two players in each row. Then, in any Nash equilibrium, the asset passes, by means of a sequence of trades, from the CGS to the CGB – in other words, the market is efficient – and in each equilibrium trade the transaction price is equal to v (except for the first row, where the transaction price is αv).

3.2 Experimental design and procedures

All the experimental sessions were conducted at the Center for Experimental Social Science (C.E.S.S.) at New York University (NYU). The subjects were recruited from the undergraduate student body of the Collage of Arts and Sciences at NYU. Subjects read the instructions silently (reproduced in Section 6), after which the instructions were read aloud by one of the experiment administrators. Subjects were invited to ask questions during the verbal instruction period. No subject reported difficulty understanding the procedures or using the computer interface. Each experimental session lasted a little more than one hour. A \$5 participation fee and subsequent earnings, which averaged about \$20, were paid in private at the end of the session.

Each experimental session consisted of 30 independent trading periods. The network structure and the trading protocol were held constant throughout a given experimental session. At the beginning of each session, each subject was randomly and independently assigned to one row of the network. This determined his type, which remained constant throughout the experiment. At the beginning of each trading period, the computer would randomly form networks by assigning subjects to the various nodes in the network. Top row subjects were assigned to top row nodes, middle row types to middle row nodes, and so on, but the assignments were otherwise random and subjects had an equal probability of being selected for each node and network.

Subjects were informed of the network structure, the trading protocol, and the payoff function. At the beginning of a trading period, they would be informed of their position in the network to which they were assigned and then would be asked to choose a bid price (the price at which they were willing to buy one unit of an asset) and an ask price (the price at which they were willing to sell one unit of the asset). Prices were denominated in terms of tokens which would be converted into dollars at the end of the experiment. Each subject had an initial endowment of 100 tokens and was allowed to choose any number (including decimals) between 0 and 100 as a bid or ask price. Subjects knew the asking price of the CGS (0 tokens) and the bid price of the CGB (100 tokens).

The computer program dialog window is shown in Section 6. The main features of the computer interface are: the large window at the left of the screen, which displays the network and the price and trading information; the View Results button in the top right corner of the screen, which allows subjects to recall the price and trading information from any previous trading period; the Bid and Ask fields at the right of the screen, where subjects enter the prices at which they are willing to buy and sell; and the message window in the lower left corner of the screen. In each period, subjects are required to enter their bids and asks in the respective fields and click the Submit button. After all subjects entered a bid and ask price, the computer executed the feasible trades according to the trading protocol. Then they saw the results of the trade in their network. After they have all clicked the OK button, the next trading period begins.

At the end of the experiment, the computer selected one trading period at random, where each period had an equal probability of being chosen, and the subject was paid an amount based on the number of tokens earned in that period. The subject's earnings in the chosen trading period would equal the initial endowment of 100 tokens, plus the trading profits (positive or negative) defined as the difference between the revenue from selling the asset (zero if the asset was not sold) and the cost of buying the asset (zero if the asset was not purchased). The payoff was defined by the formula

$$\text{payoff} = 10 + \max\{0, \ \text{trading profit}\} \tag{1}$$

Payoffs were calculated in terms of tokens and then converted into dollars.

The payoff function (1), which is used in most experimental treatments, represents a compromise between two conflicting objectives. On the one hand, we would like subjects to be rewarded on the basis of (positive or negative) trading profits, just like market traders. On the other hand, if we allow traders to keep their entire earnings, they have a strong incentive to remain passive, since they can earn an amount equal to their endowment without trading at all. The only way to remove the incentive for passivity is to subtract most of the endowment from their earnings before calculating their compensation, but this has the effect of removing the possibility of substantial losses too. As with most compromises, this one leaves us dissatisfied, but we have experimented with a payoff function that allows for more losses and find that the behavior is not too different.

Our experimental design consists of a baseline treatment followed by a number of variations to test the sensitivity of our results with respect to the degree of competition, the pricing rule, the payoff function, and finally the effect of asymmetry in the network architecture. The **baseline** treatment (**B**) uses the 2×3 and 3×3 networks, the mean-price rule ($\alpha = 0.5$), and the payoff function (1). The **auction** treatment (**A**) is identical to the baseline treatment except that it only uses the 1×3 network. Within the auction and baseline treatments, only the number of rows in the network is varied. Adding rows increases the amount of intermediation required to capture the surplus available, which allows us to test the responsiveness of pricing behavior to trading uncertainty.

The competition treatment (**C**) is the same as the baseline treatment except that it uses the 2×2 and 3×2 networks. In respect, these networks allow us to test the robustness of the results of the 2×3 and 3×3 networks in the baseline treatment to a reduction in the amount of competition in each network, where we identify competition with the number of traders in each row. The bid-price treatment (**P**) uses the 2×3 and 3×2 networks to test the robustness of the results of the baseline and competition treatments to a change in the definition of the transactions price by setting it equal to the successful bid ($\alpha = 1$). Further, the asymmetric treatment (**S**), uses a simple asymmetric 2×2 network to explore the effects of asymmetry on pricing behavior. This network is identical to the symmetric 2×2 network used in the competition treatment except that one of the nodes in the top row is connected to only one of the nodes in the bottom row.

The loss treatment (**L**) further tests the robustness of the results by using the 3×2 network and a payoff function that deducted more trading losses from subjects' earnings. More specifically, compare the payoff function (1), we increase the constant term and reducing the lower bound on trading profits such that

$$\text{payoff} \;=\; 50 + \max\{-40,\ \text{trading profit}\}\,. \tag{2}$$

Hence, if a subject makes a trading loss of more than 40, his payoff will be equal to 10 tokens. If the trading profit is non-negative, his payoff will be at least 50, that is, 40 tokens more than under payoff function (1). Other things being equal, under payoff function (2), payoffs are higher, but the incentive to trade may be smaller. The fact that the subject can now earn 50 tokens for sure by not trading, creates a significant risk of loss from trading. Figure 2 and the diagram below summarizes the experimental design. The entries of the form $a/b/c$ represent the number of networks, the number of subjects, and the number of observations per row.

Exp.	Networks	# of obs.
B	2×3	6/54/270
	3×3	6/54/180
A	1×3	5/15/150
C	2×2	9/36/270
	2×3	6/36/180
P	2×3	6/36/180
	3×2	6/36/180
S	2×2	8/32/240
L	3×2	6/36/180

[Figure 2 here]

4 Experimental Results

In this section, we present our experimental results. The novel feature of our design is the presence of intermediation. We first explore the effect of intermediation, measured by the number of rows in the network, on prices in the baseline treatment (**B**). Then, we consider the robustness of the results to changes in the amount of competition (**C**), measured by the number of columns in the network, the pricing rule (**P**), the payoff function (**L**), and the asymmetry of the network architecture (**S**). Finally, we compare the levels of efficiency, measured as the fraction of completed trades, across treatments and networks. The aim of the analysis is provides us with an insight into how experimental networks behave, as well as to test the usefulness of the theory for interpreting behavior in the laboratory.

4.1 Data description

We begin by providing an overview of some important features of the experimental data, which we summarize by reporting average bid, ask and transaction prices in a number of ways. Each experiment consists of 30 trading periods. To economize on space and to facilitate comparison across treatments and networks, instead of showing the data from each trading period, we have grouped the trading periods into terciles, corresponding to early periods (1–10), intermediate periods (11–20) and late periods (21–30). Table 1 presents the data for each tercile as a separate sub-panel. Each entry is the mean and standard deviation over the tercile, the treatment and the row from the $m \times n$ network used in the treatment.

[Table 1 here]

Table 1A shows the average *winning* bids and Table 1B shows the average *maximum* bids. The difference between the maximum bid and the winning bid occurs because there is no winning bid in the case where no trade occurs (the maximum bid is less than the corresponding ask). Similarly, Table 1C shows the average *winning* asks and Table 1D shows the average *maximum* asks. The winning ask is the asking price of a subject who actually has the asset and is successful in trading. The maximum ask, by contrast, is the maximum asking price, whether or not a trade actually occurred. Additionally, Table 1E shows the average transaction price. The transaction price corresponding to row i is the actual amount paid for the asset by the subject in row i. If no trade occurs, the transaction price is not defined and is not included in the average. A cursory examination of the data indicates two broad facts about the behavior observed in the laboratory: First, although the difficulty of solving the problem of trading in networks is sometimes massive, prices generally converge to the competitive equilibrium prices. And, secondly, trade is very efficient, in the sense that the asset typically reaches the CGB and the surplus is realized.

4.2 Prices

The auction treatment (**A**) uses 2×3 network and the baseline treatment (**B**) uses the 2×3 and 3×3 networks. Since there are three bidders in each row, any equilibrium of the trading game is efficient and the price at which the asset is traded in equilibrium is equal to 50 in the top row (the ask of the CGS is fixed at 0) and 100 in other rows (see Appendix I for details). In an experimental setting, there are many reasons why we do not at first observe the equilibrium transaction prices. Perhaps the most important reason is strategic uncertainty: a subject bidding for an asset has little information about the price at which he can re-sell the asset, unless he happens to be in the bottom row and can sell the asset to the CGB for the price of 100.

4.2.1 Convergence

Uncertainty about the resale price may cause subjects to shave their bids to protect themselves against the possibility of selling at a loss, or failing to sell at all. Although subjects are randomly matched with different subjects each trading period, as the game is repeated, they can learn from experience and their uncertainty diminishes. As a result, competition may be expected to increase and eventually cause convergence of the actual price to its equilibrium level. This is exactly what we observe in the baseline treatment. Result 1 summarizes the effect of repetition on the average winning bid, ask and transaction prices.

Result 1 (convergence) *In the baseline treatment, the average winning bid and ask prices are initially far below their equilibrium values, but they converge rapidly after several repetitions. Transaction prices, being the average of winning bids and asks, also converge rapidly to the equilibrium values, with the exception of the bottom row of the* 3×3 *network, where the average transaction price actually falls between the second and third terciles.*

The relevant support for Result 1 comes from Table 1. The average winning bids converge especially fast. In the second and third terciles, all bids are all within one percent of the equilibrium bid of 100. The average winning asks converge more slowly, but by the third tercile they are within one percent of 95 in the top row and within one percent of the equilibrium ask of 100 in the bottom row of each network. Only in the middle row of the 3×3 network there is a curious drop in the winning asks in the third tercile. Overall, given subjects' uncertainty about the possibility of reselling the asset, it take remarkably little time to prices to converge. The maximum bids and asks should show similar patterns (in practice, the patterns are almost identical once we aggregate by tercile). Next, we look more closely at the effect of the network architecture on pricing by comparing behavior in corresponding rows across network and in different rows within a given network.

4.2.2 Intermediation

We first examine the pricing behavior of subjects belonging to corresponding rows across network. The 1×3 network used in the auction treatment (**A**) and the 2×3 and 3×3 networks used in the baseline treatment (**B**) differ only in the number of rows. The more rows, the greater the amount of intermediation required to transfer the asset from the CGS to the CGB. One might expect that more intermediation would reduce the rate of convergence to the equilibrium price, but that does not appear to be the case. In fact, pricing behavior is quite robust to variation in the number of rows in the network. Result 2 summarizes the behavioral regularities in this regard by comparing average winning bids and asks in rows that have similar positions relative to the CGB.

Result 2 (intermediation) *The rates of convergence of the average winning bids and asks to the equilibrium value are not significantly different in the corresponding rows of the three networks in the baseline and auction treatments, that is, the* 1×3 *network and the bottom rows of the* 2×3 *and* 3×3 *networks, and the top and middle rows of the* 2×3 *and* 3×3 *networks, respectively.*

The support for Result 2 comes again from Table 1. We reorganize this information below. For the bid prices (top panel), the only case where the rate of convergence differs is during the first tercile, where the top row of the 2×3 network has significantly lower bids than in the middle row of the 3×3 network. For the ask prices (bottom panel), there is some variation in the first tercile, but in later terciles the corresponding rows have similar prices, except for the odd drop in asking prices in the middle row of the 3×3 network in the last tercile. Apart from this small difference, the asks in corresponding rows do not differ significantly across networks after the first tercile. Note that we do not include the asks in the bottom rows in each network, since subjects very quickly realized that they could ask for 100 from the CGB. Figure 3 below presents, turn by turn, the data on the average winning bids.

Average wining bids

Network	Row	1	2	3
1×3	1	95.9	100	100
2×3	2	90.0	99.9	100
3×3	3	97.2	100	100
2×3	1	77.4	99.9	100
3×3	2	88.2	99.9	100

Average wining asks

Network	Row	1	2	3
2×3	1	70.7	87.8	94.3
3×3	2	80.1	90.8	80.3

[Figure 3 here]

4.2.3 Spreads

Another interesting feature of the network architecture is the price spread between different rows of a given network. Again, the critical factor is the uncertainty about resale as measured by the distance from the CGB. In the auction treatment (**A**), which uses the 1×3 network, the successful buyer knows that he can always sell the asset to the CGB for the price of 100. Thus, the absence of uncertainty guarantees aggressive bidding in line with the predictions of equilibrium. By contrast, in the 2×3 and 3×3 network used in the baseline treatment (**B**), subjects cannot be sure of the price at which they can resell the asset. This strategic uncertainty would tend to depress the bids and asks, and suggest that, in the 3×3 network for example, the transaction prices between first row seller and the second row buyer will be lower than the transaction prices between the second row seller and the third row buyer. Our next result provides information about the evolution of transaction prices within each network in the baseline treatment.

Result 3 (spreads) *The average transaction prices in a given network are increasing in the row index for each network in the baseline treatment and in each tercile, with the exception of the 3×3 network, where in the third tercile the average transaction price is lower in the bottom row than in the middle row.*

Evidence for Result 4 is also provided by Table 1. We present the relevant data from Table 1 below by comparing the average transaction prices across rows within a given network and tercile. Recall that the pricing rule in the baseline treatment ($\alpha = 0.5$) and the fact that the CGS always asks 0 together imply that the cost of the asset to the winning bidder in the first row is never more than 50. This makes it hard to compare bids in the first row with bids in the subsequent rows. For this reason we prefer to compare transaction prices between rows. Furthermore, because the bidders in the second row pay a price that depends on the seller's ask as well as on their bids, they may be inclined to keep their bids lower than if they were buying from the CGS. In any case, the first row subjects can only make a profit if they sell the asset for more than they paid for it and this suggests that the transaction price between the CGS and the first row will be lower than the transaction price between the first row and the second row.

Average transaction prices

Network	Row	1	2	3
2×3	1	38.7	49.9	50.0
2×3	2	80.3	93.9	97.2
3×3	1	38.8	49.9	50.0
3×3	2	81.0	92.9	97.7
3×3	3	88.6	95.4	90.2

Summarizing, the data from the auction (**A**) and baseline (**B**) treatments show that the experimental design is stable and easy to understand, allowing us to test the basic elements of the theory, most importantly, whether observed behavior corresponds to the equilibrium predictions of the theory. Overall, it appears that there are strong forces leading subjects to the equilibrium of the game. Further, the convergence comes from below, that is, subjects begin by bidding and asking low prices and gradually raise their prices as they become more confident about the behavior of other subjects and as the bid and ask prices convergence to equilibrium level.

4.3 Sensitivity

In what follows, we consider a number of variations on the baseline treatment with a view to testing the robustness of the results to different aspects of the experimental design.

4.3.1 Competition

In the 2×3 and 3×3 networks used in the baseline treatment (**B**), there are three bidders ($n = 3$) in each row. In theory, Bertrand competition will guarantee an equilibrium price of 100 as long as there are at least two bidders in each row. In the laboratory, we do not necessarily expect perfectly competitive behavior when the number of bidders is small. It is therefore of particular interest to see whether the results of the baseline treatment hold up when we reduce the number of subjects in each row. To this end, in the competition treatment (**C**), we reduce the number of bidders in each row from three to two ($n = 2$), but keep all other parameters the same as in the baseline treatment. Specifically, we compare, row by row, the 2×3 and 3×3 networks used in the the baseline treatment with the 2×2 and 3×2 networks used in the competition treatment, respectively. Our next results report that, within a given row, in the 2×3 and 2×2 networks there is no significant difference between the pricing behaviors, but the situation clearly reverses, particularly in early periods, in the 3×3 and 3×2 networks.

Result 4 (competition) *Comparing networks with low ($n = 2$) and high ($n = 3$) competition, the rates of convergence of the average winning bids and asks are not significantly different in the 2×3 and 2×2 networks. In the 3×3 and 3×2 networks, the differences are larger. As a result, the price spreads between transaction prices are more noticeable in the 3×2 network, especially in the top and middle rows.*

Support for Result 4 is also based on the data from Table 1. Below, we compare, row by row, the data from Table 1 on the average winning bids (top panel) and asks (middle panel), and the transaction prices (bottom panel) in the baseline and competition treatments. The 2×3 and 2×2 networks have very similar bidding and asking behavior. The situation is more complex when there is more intermediation. In the 3×3 and 3×2 networks, there are interesting and significant differences in the average winning bids and asks in the top and middle rows, but not in the bottom row. As a result, average transaction prices show the same patterns when competition is high ($n = 3$) as when it is low ($n = 2$), though the levels are different, in particular, in the 3×3 and 3×2 networks.

An exception occurs in the transaction price between the middle and bottom rows of the 3×3 network in the third tercile. In this case, because of the curious drop in the winning asks in the middle row, the transaction price is surprisingly low, lower than the transaction price between the same rows in the 3×2 network. With this exception, the general pattern is that competition increases the transaction price in any given row. That is, the rates of convergence are slower in the 3×2 network than in the 3×3 network so less competition does make a difference here. Nevertheless, behavior reach the neighborhood of equilibrium at the end of the experiment also when competition is low. Figure 4 below presents, in graphical form, the data on the average winning bids.

Average wining bids

Networks	Row	1	2	3
2×3 vs. 2×2	1	$77.4 - 70.9$	$99.9 - 94.4$	$100 - 99.9$
2×3 vs. 2×2	2	$90.0 - 89.9$	$99.9 - 96.9$	$100 - 98.4$
3×3 vs. 3×2	1	$77.7 - 36.6$	$99.9 - 65.4$	$100 - 91.7$
3×3 vs. 3×2	2	$88.2 - 58.5$	$99.9 - 78.5$	$100 - 91.5$
3×3 vs. 3×2	3	$97.2 - 84.4$	$100 - 96.7$	$100 - 98.9$

Average wining asks

Networks	Row	1	2	3
2×3 vs. 2×2	1	70.7 – 76.5	87.8 – 85.1	94.3 – 85.8
2×3 vs. 2×2	2	94.5 – 97.5	97.0 – 99.9	99.8 – 100
3×3 vs. 3×2	1	73.9 – 50.3	85.9 – 71.5	95.4 – 87.4
3×3 vs. 3×2	2	80.1 – 70.8	90.8 – 88.8	80.3 – 96.3
3×3 vs. 3×2	3	99.9 – 97.3	100 – 99.8	98.2 – 99.9

Average transaction prices

Networks	Row	1	2	3
2×3 vs. 2×2	1	38.7 – 35.4	47.2 – 47.4	50.0 – 50.0
2×3 vs. 2×2	2	80.3 – 83.2	93.9 – 91.0	97.2 – 92.1
3×3 vs. 3×2	1	38.8 – 18.3	49.9 – 32.7	50.0 – 45.9
3×3 vs. 3×2	2	81.0 – 54.4	92.9 – 75.0	97.7 – 89.5
3×3 vs. 3×2	3	88.6 – 77.6	95.4 – 92.8	90.2 – 97.6

[Figure 4 here]

4.3.2 Pricing rule

In the baseline (**B**) and competition (**C**) treatments, the price paid for the asset is the average of the bid and the ask ($\alpha = 0.5$), and the corresponding amount of tokens is transferred from the buyer to the seller. The bid-price treatment (**P**), which uses the 2×3 and 3×2 networks, differs from the same networks in the baseline and competition treatments only in setting the transaction price equal to the bid price ($\alpha = 1$). Intuitively, this could slow convergence to equilibrium by making subjects less willing to bid aggressively. Nevertheless, our next result reports that the rates of convergence are quite similar under bid-price and average pricing rules. Since in the bid-price treatment we set the transaction price equal to the successful bid, we restrict attention to the evolution of the average winning bids. Below, we reorganize the evidence for the result, which is also presented in Figure 5 below.

Result 5 (pricing rule) *In the 2×3 network, there are no significant differences between the rates of convergence to equilibrium bid prices in the bid-price and baseline treatments, especially in the second and third terciles. In the 3×2 network, the rates of convergence in the competition treatment are slightly higher than in the bid-price treatment,* *except in the top row in the third tercile where the difference is more significant.*

Average wining bids

Treatments	Row	1	2	3
B vs. **P**	1	77.4 – 89.4	99.9 – 98.5	100 – 99.4
B vs. **P**	2	90.0 – 95.1	99.9 – 98.8	100 – 99.5
C vs. **P**	1	36.6 – 35.8	65.4 – 57.8	91.7 – 70.0
C vs. **P**	2	58.5 – 62.1	78.5 – 76.8	91.5 – 85.2
C vs. **P**	3	84.4 – 79.2	96.7 – 86.6	98.9 – 91.4

[Figure 5 here]

4.3.3 Payoffs

The particular payoff function used in the Auction (**A**), baseline (**B**), competition (**C**) and bid-price (**P**) treatments discussed above excludes the possibility of substantial trading losses. The loss treatment (**L**) uses a payoff function that deducts more trading losses from subjects' earnings than in the other treatments. Consequently, it seems plausible to expect that price convergence will be slower than in the baseline treatment or fail to occur completely. Our next result confirms this conjecture and shows that losses have the most significant effect on convergence, because they make subjects less willing to bid aggressively for the asset. We can thus conclude that price convergence can be achieved upon repetition, but not in all market environments.

The evidence for the result is given below by comparing the average winning bids (top panel), asks (middle panel) and transaction prices (bottom panel) in the loss and competition (**C**) treatments, using the 3×2 network. In each case, the prices corresponding to the loss treatment are lower and the gap between the two treatments often widens over time. Furthermore, in the loss treatment, convergence to the equilibrium price has not occurred by the end of the experiment and it is not clear that further repetitions would lead to complete convergence. Also note that although the differences in the third row are smaller than in the other rows, except for an upturn toward the end of the experiment, there is no clear indication of convergence to the equilibrium price in the loss treatment. Finally, Figure 6 also presents, turn by turn, the data on the average winning bids. Summarizing,

Result 6 (payoffs) *In all three rows of the 3×2 network, the rates of convergence in the loss treatment are much slower than in the competition treatment and in some cases fails to converge.*

Average wining bids

Treatments	Row	1	2	3
C vs. **L**	1	36.6 – 26.9	65.4 – 38.8	91.7 – 51.3
C vs. **L**	2	58.5 – 53.3	78.5 – 62.2	91.5 – 70.0
C vs. **L**	3	84.4 – 70.0	96.7 – 73.9	98.9 – 80.0

Average wining asks

Networks	Row	1	2	3
C vs. **L**	1	50.3 – 43.7	71.5 – 53.1	87.4 – 60.8
C vs. **L**	2	70.8 – 62.9	88.8 – 68.2	96.3 – 73.1
C vs. **L**	3	97.3 – 94.7	99.8 – 98.1	99.9 – 95.9

Average transaction prices

Networks	Row	1	2	3
C vs. **L**	1	18.3 – 13.4	32.7 – 19.4	45.9 – 25.6
C vs. **L**	2	54.4 – 78.5	75.0 – 57.7	89.5 – 65.4
C vs. **L**	3	77.6 – 66.5	92.8 – 71.0	97.6 – 76.6

[Figure 6 here]

4.3.4 Asymmetry

So far we have looked at networks in which each row is completely connected to the adjacent rows, that is, all possible links are assumed to be present. These networks have the advantage of symmetry, that is, the nodes in a given row are essentially identical. This symmetry allows us to pool the data generated by the subjects in a given row. Nevertheless, there are many other architectures we can study and this will be an interesting extension of our current work. Most importantly, there are many interesting questions that can only be answered using an asymmetric network architecture, such as what are the differences in price setting behavior among subjects who have different numbers of trading partners.

The asymmetric network we included is very simple, but it gives some hints as to the kinds of phenomena that might be found in more complex asymmetric networks. In one respect, our treatment does not fully exploit the asymmetry of the network when it comes to price setting because, in order to make the computer program easier to use, subjects could only choose a single bid and ask price for all their potential trading partners. In a symmetric network this may not be very restrictive since, from the point of view of a single subject, the buyers and sellers are essentially the same. In an asymmetric network, the assumption is restrictive, because a buyer may want to offer different prices to different potential sellers depending on the amount of competition the buyer faces.

For simplicity, the asymmetric treatment (**S**) uses the 2×2 network, which was also used in the competition treatment (**C**). The symmetric (left panel) and asymmetric (right panel) 2×2 networks are illustrated in Figure 7 below. We label the top row nodes 11 and 12 and the bottom row nodes 21 and 22 and assume that in the asymmetric network 11 is connected to both 21 and 22, whereas 12 is connected only to 22. Thus, 22 is a monopsonist if 12 has the asset, and a duopsonist if 11 has it. Then it seems likely that 22 will exploit his "bargaining power" by offering lower bids, even though he has to offer the same price to 11 and 12, and he may lose the asset to 21 if his bid is too low. But against this he weighs the extra profits he will get if 12 has the asset.

[Figure 7 here]

Our next results report that what appears to happen is that the price-setting behavior of the subjects is symmetric in each row. The effect of 22's "bargaining power" is simply to lower *all* prices compared to the symmetric network. The support for the result comes from Table 2 below. Table 2 below provides a first indication by summarizing the average bids and asks for the different nodes in the asymmetric 2×2 network by tercile. As claimed above, we see that the behavior of 11 and 12 is very similar, as is the behavior of 21 and 22. In addition, Table 3 shows the average winning bids, asks and transaction prices for different trading paths in the asymmetric network (for trades between 11 and 21, between 11 and 22, and between 12 and 22) and compares the results to the data from the asymmetric network. Again we find the prices are generally independent of the path taken. Thus, the effect of asymmetry on prices tend to propagate through the network as a change in prices in one part of the network affects what traders are prepared to bid and ask elsewhere. There also appear to be some asymmetries in the proportion of trades corresponding to each of these paths, but given the similarity of the pricing behavior, it is hard to ascribe any clear meaning to these asymmetries, even assuming they are significant. Concluding,

Result 7 (asymmetry) *The pricing behavior of subjects in a given row in the asymmetric 2×2 network is very symmetric. The effect of the network asymmetry is revealed by generally lower prices and slower convergence compare to the symmetric 2×2 network.*

[Table 2 here]
[Table 3 here]

4.4 Efficiency

The uncertainty of trade in networks provides an important source of market imperfection. Thus, the efficiency of trade is one of the main concerns in the study trading in networks. The goal is to identify how the network architecture influences the efficiency of trade. The greater the incompleteness of the network, the more intermediation is required to achieve an efficient outcome. The cost and uncertainty of intermediation provide an important source of market imperfection. Recall that trade between two rows requires that at least one bid is higher than or equal to the seller's ask. Thus, strategic uncertainty (about what other subjects will do) inevitably requires a period of learning and during this period trades may not be completed. Further, even later in the experiment trade may break down if subjects make mistakes about the prices that are likely to be bid or asked by their opponents. Clearly, the transaction prices only effect the distribution of the surplus among traders. Efficiency depends only on whether the asset reaches the CGB and the surplus is realized. The efficiency of trade is reported in Table 4, which displays, row by row, the fraction of completed trades in each tercile of periods for each treatment and network. The first row is excluded because there is no possibility of incomplete trades.

[Table 4 here]

On the whole, we are pleasantly surprised by the extent of trading. In the baseline treatment (**B**), the level of efficiency is generally higher in the 2×3 network than in the 3×3 network (the degree of intermediation is higher). Also, the level of efficiency increases significantly through the three terciles in the 3×3 network, but is quite high and essentially flat in the 2×3 network. This suggests that it takes subjects longer to learn to coordinate when there is more intermediation. Reducing the amount of competition by reducing the number of bidders in each row reduces efficiency, but not by much. Comparing the 2×3 network from the baseline treatment with the 2×2 network from the competition treatment (**C**), we see that efficiency is lower in the 2×2 network in the first and second tercile, but the difference is not significant by the last tercile. Similarly, comparing the 3×3 network with the 3×2 network, we see that efficiency is lower in the 3×2 network in each tercile, though the differences are quite small. Thus, we conclude that more intermediation or less competition can lower the efficiency of trade.

Changing the transaction price from the average of the winning bid and ask to the winning bid price reduces efficiency in both the 2×3 and 3×2 network used in the bid-price treatment (**P**) compare to the corresponding networks in the baseline and competition treatments. One of the interesting features of the bid-price treatment is that efficiency declines slightly from the first to the third tercile in the 2×3 network, whereas in the 3×2 network efficiency increases sharply from the first to the second tercile and increases modestly between the second and third terciles. In the loss treatment (**L**), efficiency is lower, as one would expect. In the 3×2 network used in the competition treatment, efficiency is steadily increasing over time, whereas in the same network in the loss treatment, efficiency increases and then decreases. In the first tercile, the efficiency in the loss treatment is higher, but in the second and third treatment it is lower than in the competition treatment. Finally, efficiency is much lower in the first tercile in the 2×2 network used in the asymmetric treatment (**S**) than in the same network used in the competition treatment, but the gap narrows in the second and third terciles. Our last result summarizes this discussion.

Result 8 (efficiency) *On the whole, efficiency of trade is very high. The levels of efficiency appear to be lower when there is more intermediation or less competition. Further, that trading rules are important for efficiency: when the transaction price equal to the bid price, subjects experience trading losses or asymmetric trading links, the level of efficiency appears to be lower.*

5 Conclusion

In this paper we examine the effects of intermediation and competition on different properties of the market behavior. So far we have almost exclusively looked at rectangular arrays in which each row is completely connected to the adjacent rows, that is, all possible links are assumed to be present. There are many other architectures we can study and this will be an interesting extension of the current paper. In addition, there are many other questions that can be addressed using this design or variations thereof. An important class of phenomena requires us to introduce randomness. Here we mention three possibilities.

- *Random endowments.* A trader's endowment places an upper limit on what he can bid for an asset and serves as a "liquid constraint." By making endowments random we introduce liquidity shocks which will change pricing both directly, by constraining bids, and indirectly by reducing competition for bidders and lowering resale prices for intermediaries.

- *Random graphs.* Random graphs are intrinsically interesting because they introduce uncertainty about the availability of counterparties to trade with. They also give rise to interesting strategic phenomena. For example, if the number of bidders in an auction is uncertain and with positive probability the number of bidders is one, the only equilibrium involves mixed strategies. Further, randomness can propagate through the network as a change in prices in one part of the network affects what traders are prepared to bid and ask elsewhere.

- *Random values.* Uncertainty about the values assigned to the asset by the CGS and the CGB introduces uncertainty about the probability of trade and the possibility of learning the value of the asset over time. Our simple framework can provide insight into how these important phenomena will be affected by network architectures.

Furthermore, there is a vast number of interesting network architectures. While the small networks we studied are very insightful, especially in experimental contexts, the development of the theory depends on properties of networks that can be generalized. In order to determine which factors are important in explaining market behavior, it will be necessary to investigate a large class of networks in the laboratory. Fortunately, our novel experimental design, employing graphical representations of networks of traders, enables us to do this systematically and efficiently. More experiments can provide us a great opportunity to test the predictions of the theory and, at the same time, evaluate and develop our experimental methodology. In addition to testing theories, we also hope to study the effects of variables about which our existing theory has little to say.

6 Sample Instructions (Baseline 2×3)

This is an experiment in the economics of decision-making. Research foundations have provided funds for conducting this research. Your payoffs will depend on your decisions and the decisions of the other participants, as well as on chance. If you follow the instructions and make careful decisions, you may earn a considerable amount of money.

The entire experiment should be complete within an hour and a half. At the end of the experiment you will be paid privately. At this point, take a minute to write down the number of the computer you are using as it appears on the top of the monitor. At the end of the experiment, you will use your computer number to claim your earnings.

At this time, you will receive $5 as a participation fee. Details of how you will make decisions will be provided below. During the experiment we will speak in terms of experimental tokens instead of dollars. Your payoffs will be calculated in terms of tokens and then translated at the end of the experiment into dollars at the following rate:

$$1 \text{ Token} = 1 \text{ Dollar}$$

The experiment is divided into 30 independent and identical trading periods. In each period, you will be asked to submit bids (prices at which you are willing to buy) and asks (prices at which you are willing to sell) for a single unit of an indivisible asset. Trades take place in a set of interconnected markets represented by a six-person network. You will only be able to trade with participants to whom you are connected in this network.

The experiment starts by having the computer randomly assign each subject to one of two rows: top or bottom. You have an equal probability of being assigned to each row and your row assignment will remain unchanged throughout the experiment. Before the start of each period, you will be randomly assigned to one of the positions in one of the networks. The positions are labeled with the letters A through F. The top row consists of positions (A, B, C), and the bottom row consists of positions (D, E, F).

Each period starts by having the computer randomly form six-person networks by selecting one participant of type-A, one of type-B, one of type-C, and so on. If you were initially designated a top row player, you will be assigned to a top row position in one of the networks, and similarly if you are a bottom row player. Your type (A, B, C, D, E, F) will be displayed in the top right hand corner of the program dialog window (see attachment 1).

[Attachment 1 here]

The networks formed in each period depend solely upon chance and are independent of the networks formed in any of the other periods. That is, in any network each top-row participant is equally likely to be chosen as type-A participant for that network, and similarly with participants of types B and C. Likewise, in any network each bottom-row participant is equally likely to be chosen as type-D participant for that network, and similarly with participants of types E and F.

Note again that your row assignment will remain unchanged throughout the experiment but your type and network may change from period to period. In each period, it depends solely on chance.

The network is displayed in the large window that appears in the center of the program dialog window (see attachment 1). A line segment between any two types indicates that they are connected and, hence, are allowed to trade. The arrowhead points from the seller to the buyer. In the network used in this experiment, each of the types in the top row (A, B, C) can trade with each of the types in the bottom row (D, E, F) and vice versa.

The asset is initially held by a computer-generated seller. The computer-generated seller is always willing to sell one unit of the asset for a price of zero tokens. In addition, there is a computer-generated buyer who is always willing to buy one unit of the asset at a price of 100 tokens. The computer-generated seller can only sell the unit to the types in the top row (A, B, C). The computer-generated buyer can only buy the unit from the types in the bottom row (D, E, F). Note that the computer-generated seller and buyer do not appear in the network displayed in the program dialog window (see attachment 1).

A trading period

Next, we will describe in detail the process that will be repeated in all 30 periods and the user interface that you will use to make your decisions. Each period starts by having the computer randomly form six-person networks by selecting one participant of each type (A, B, C, D, E, F). At the start of each period, each participant receives an endowment of 100 tokens. You will use these tokens to pay for the asset when you buy and will receive these tokens in exchange for the asset when you sell. The trading protocol is defined by the following rules.

All trades must move the asset “downward”:

- The types in the top row can only buy from the computer-generated seller and can only sell to the bottom row. For example, type A can buy from the computer generated seller and can sell to types D, E or F.

- The types in the bottom row can only buy from the top row and sell to the computer-generated buyer. For example, type D can buy from types A, B or C and sell to the computer-generated seller.

In each period, you will be asked to submit a single bid and a single ask to the sellers and buyers with whom you are allowed to trade:

- You will submit a single bid to the sellers to whom you are connected by the network, indicating the price at which you are willing to buy one unit of the asset.

- You will submit a single ask to the buyers to whom you are connected by the network, indicating the price at which you are willing to sell one unit of the asset.

When you are ready to make your decisions, use the mouse to position the cursor in the Bid Input field on the right of the dialog window (see attachment 1) and use the keyboard to enter the number (including decimals) of tokens between 0 and 100 that you wish to bid. You enter a price in the Ask Input field the same way. Once you have entered the bid and ask, confirm your decisions by clicking the Submit button. Once you have clicked the Submit button, your decisions cannot be revised.

When everyone has submitted their sealed bid and ask, you will observe the bids and asks of all other participants, the actual prices at which the asset was traded and the sequence of trades. This information is displayed in the large window that appears in the center of the dialog window (see attachment 2). Bids and asks are colored blue, and the actual prices at which the asset was traded are colored green. The letter to the left of each price indicates whether it is an ask (A) or a bid (B). Each bid is displayed above the type of participant who has submitted this bid, and each ask is displayed below the type of participant who has submitted this ask. The asks of the computer-generated seller and bids of the computer-generated buyer are indicated by CA and CB, respectively.

[Attachment 2 here]

Trades are executed sequentially. First, trades between the computer-generated seller and the buyers in the top row take place, followed by trades between the seller in the top row and the buyers in the bottom row, and followed by trades between the seller in the bottom row and the computer-generated buyer. At each stage, a trade occurs only if a buyer has submitted a bid that is at least as high as the seller's ask. If there is more than one bid that is greater than or equal to the asking price, the asset is transferred from the seller to the buyer with the highest bid. If two buyers tie for the highest bid, the asset will be assigned to one of the buyers at random. The buyer pays the seller the number of tokens equal to the average of the bid and ask. Trading stops at any stage where no buyer bids as much as the seller's ask. In that case, the asset remains with the seller.

This completes the first of 30 trading periods. To move on to the second period, press the OK button on the bottom right hand corner of the program dialog window. Note that after one minute the program will move automatically to the second period, but you will always be able to review the results of this period later in the experiment by choosing it and clicking on the View Results button on the top right hand corner of the program dialog window (see attachment 2). After letting you observe the results of the first period, the second period will start by having the computer randomly forming new groups of participants in networks.

This process will be repeated until all the 30 independent and identical trading periods are completed. Throughout the experiment please pay careful attention to the messages window at the bottom of the program dialog window (see attachment 1). At the end of the last round, you will be informed the experiment has ended.

Payoffs

Your trading profit in each period can be summarized by the formula:

Profit = (selling price) - (buying price)

The buying price is the actual price you paid for the asset if you traded and zero otherwise. The selling price is the actual price you received for the asset if you sold the asset and zero if you did not trade or if you bought and did not sell. Your total earnings in each trading period are equal to your initial endowment of 100 tokens plus your trading profits, positive or negative.

Your final payoff in the experiment is determined as follows. At the end of the experiment, the computer will randomly select one period in which to execute the trades "for real". The period selected depends solely on chance.

If the number of tokens you earned in that period is less than 100, you will receive 10 tokens to keep. If the number of tokens you earned in that period is at least 100, you will receive that amount minus 90 tokens to keep. At the end of the experiment, the tokens will be converted into money. Each token will be worth $1. You will receive your payment as you leave the experiment.

Rules

Please do not talk with anyone during the experiment. We ask everyone to remain silent until the end of the last round. Your participation in the experiment and any information about your earnings will be kept strictly confidential. Your payments receipt and participant form are the only places in which your name and social security number are recorded. If there are no further questions, you are ready to start. An instructor will approach your desk and activate your program.

References

[1] Charness, G., M. Corominas-Bosch and G. Fréchette (2005) "Bargaining and Network Structure: An Experiment." Forthcoming *Journal of Economic Theory*.

[2] Corominas-Bosch, M. (2004) "Bargaining in a Network of Buyers and Sellers." *Journal of Economic Theory*, *115*, 35-77.

[3] Forsythe, R., T. Palfrey and C. Plott (1982) "Asset Valuation in an Experimental Market," *Econometrica*, *50*, pp. 537-567.

[4] Forsythe, R., T. Palfrey and C. Plott (1984) "Futures Markets and Informational Efficiency: A Laboratory Examination," *Journal of Finance*, *39*, pp. 955-981.

[5] Gale, D. and S. Kariv (2007) "Financial Networks." Forthcoming *American Economic Review* Papers and Proceedings.

[6] Goyal, S. (2004) "A Survey of Models of Network Formation: Stability and Efficiency." In *Group Formation in Economics; Networks, Clubs and Coalitions*, ed. Gabrielle Demange and Myrna Wooders, Cambridge University Press, Cambridge U.K.

[7] Jackson, M. (2004) "A Survey of Models of Network Formation: Stability and Efficiency." In *Group Formation in Economics; Networks, Clubs and Coalitions*, ed. Gabrielle Demange and Myrna Wooders, Cambridge University Press, Cambridge U.K.

[8] Kosfeld, M. (2004) "Economic Networks in the Laboratory: A Survey." *Review of Network Economics*, *3*, pp. 20-41.

[9] Plott, C. and S. Sunder (1982) "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models," Journal of Political Economy, 90, pp. 663-698.

[10] Plott, C. and S. Sunder (1988) "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica*, *56*, pp. 1085-1118.

[11] Sunder, S. (1995) "Experimental Asset Markets: A Survey," in *Handbook of Experimental Economics*, A. Roth and John Kagel, eds., Princeton, NJ: Princeton University Press (1995), pp. 445-500.

Table 1A: Average winning bid

(by treatment, network, row, and tercile)

Periods 1-10

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	95.9	2.36	--	--	--	--
B 2×3	77.4	2.55	90.0	1.07	--	--
B 3×3	77.7	2.71	88.2	1.63	97.2	0.84
C 2×2	70.9	1.95	89.9	0.88	--	--
C 3×2	36.6	1.45	58.5	1.83	84.4	1.43
P 2×3	90.0	1.62	95.1	0.93	--	--
P 3×2	40.0	1.71	65.0	1.71	79.2	1.27
L 3×2	26.9	1.68	53.3	1.47	70.0	1.89
S 2×2	54.6	2.15	61.6	1.79	--	--

Periods 11-20

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	100.0	0.00	--	--	--	--
B 2×3	99.9	0.11	99.9	0.06	--	--
B 3×3	99.9	0.09	99.9	0.04	100.0	0.04
C 2×2	94.4	0.97	96.9	0.32	--	--
C 3×2	65.4	1.21	78.5	0.77	96.7	0.42
P 2×3	98.3	0.24	98.8	0.21	--	--
P 3×2	60.4	1.28	77.4	1.01	86.6	0.55
L 3×2	38.8	3.01	62.2	2.03	73.9	2.02
S 2×2	73.8	1.94	76.9	1.20	--	--

Periods 21-30

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	100.0	0.00	--	--	--	--
B 2×3	100.0	0.00	100.0	0.00	--	--
B 3×3	100.0	0.00	100.0	0.00	100.0	0.00
C 2×2	99.9	0.03	98.4	0.23	--	--
C 3×2	91.7	0.98	91.5	0.54	98.9	0.19
P 2×3	99.4	0.09	99.5	0.08	--	--
P 3×2	70.5	0.75	85.2	0.66	91.4	0.42
L 3×2	51.3	4.03	70.0	2.28	80.0	1.96
S 2×2	86.4	1.52	85.8	0.93	--	--

Table 1B: Average maximum bid
(by treatment, network, row, and tercile)

Periods 1-10

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	95.9	2.36	--	--	--	--
B 2×3	82.5	2.33	88.8	1.14	--	--
B 3×3	77.7	2.71	87.5	1.62	96.4	0.72
C 2×2	70.9	1.95	89.1	0.89	--	--
C 3×2	36.6	1.45	56.5	1.55	82.6	1.71
P 2×3	90.0	1.62	95.0	0.85	--	--
P 3×2	40.0	1.71	61.8	1.58	76.7	1.05
L 3×2	26.9	1.68	51.3	1.42	65.2	2.02
S 2×2	54.6	2.15	60.4	1.50	--	--

Periods 11-20

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	100.0	0.00	--	--	--	--
B 2×3	100.0	0.00	99.9	0.08	--	--
B 3×3	99.9	0.09	99.9	0.04	99.8	0.08
C 2×2	94.4	0.97	96.3	0.40	--	--
C 3×2	65.4	1.21	78.3	0.73	96.5	0.40
P 2×3	98.3	0.24	98.7	0.20	--	--
P 3×2	60.4	1.28	77.0	0.98	86.6	0.52
L 3×2	38.8	3.01	63.4	1.91	75.1	1.84
S 2×2	73.8	1.94	77.5	1.08	--	--

Periods 21-30

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	100.0	0.00	--	--	--	--
B 2×3	100.0	0.00	99.5	0.30	--	--
B 3×3	100.0	0.00	100.0	0.00	100.0	0.00
C 2×2	99.9	0.03	98.3	0.23	--	--
C 3×2	91.7	0.98	91.6	0.52	98.7	0.21
P 2×3	99.4	0.09	99.4	0.09	--	--
P 3×2	70.5	0.75	84.6	0.63	91.5	0.35
L 3×2	51.3	4.03	70.2	2.10	80.9	1.66
S 2×2	86.4	1.52	87.9	0.84	--	--

Table 1C: Average winning ask
(by treatment, network, row, and tercile)

Periods 1-10

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	98.1	1.48	--	--	--	--
B 2×3	70.7	1.41	94.5	1.51	--	--
B 3×3	73.9	2.32	80.1	4.53	99.9	0.12
C 2×2	76.5	1.55	97.5	1.00	--	--
C 3×2	50.3	1.66	70.8	2.03	97.3	1.00
P 2×3	66.8	4.39	89.9	3.89	--	--
P 3×2	50.6	1.39	66.8	1.20	83.4	2.25
L 3×2	43.7	1.73	62.9	2.03	94.7	1.39
S 2×2	47.5	1.61	92.3	1.59	--	--

Periods 11-20

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	95.0	2.92	--	--	--	--
B 2×3	87.8	1.41	97.0	1.60	--	--
B 3×3	85.9	2.62	90.8	3.63	100.0	0.00
C 2×2	85.1	2.00	99.9	0.04	--	--
C 3×2	71.5	1.21	88.8	0.59	99.8	0.07
P 2×3	70.6	4.76	84.6	5.05	--	--
P 3×2	62.2	1.88	61.9	4.32	94.8	0.80
L 3×2	53.1	2.12	68.2	2.15	98.1	0.91
S 2×2	61.4	1.60	98.6	0.49	--	--

Periods 21-30

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	94.1	2.71	--	--	--	--
B 2×3	94.3	0.84	99.8	0.16	--	--
B 3×3	95.4	1.26	80.3	5.16	98.2	1.72
C 2×2	85.8	1.22	100.0	0.03	--	--
C 3×2	87.4	0.67	96.3	0.34	99.9	0.03
P 2×3	67.4	5.72	86.3	4.85	--	--
P 3×2	70.2	2.15	71.3	3.07	97.1	0.62
L 3×2	60.8	2.01	73.1	2.27	95.9	2.86
S 2×2	65.5	1.83	99.9	0.03	--	--

Table 1D: Average maximum ask
(by treatment, network, row, and tercile)

Periods 1-10

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	99.5	0.50	--	--	--	--
B 2×3	80.2	1.27	98.3	0.61	--	--
B 3×3	82.1	2.03	93.5	1.15	100.0	0.00
C 2×2	83.3	1.15	98.6	0.79	--	--
C 3×2	54.3	1.44	71.6	1.61	96.8	1.48
P 2×3	89.2	1.48	99.3	0.52	--	--
P 3×2	61.2	1.62	72.2	1.34	92.6	1.21
L 3×2	45.7	1.46	65.3	1.58	91.5	2.13
S 2×2	60.7	1.88	94.7	1.38	--	--

Periods 11-20

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	100.0	0.00	--	--	--	--
B 2×3	97.4	0.56	100.0	0.00	--	--
B 3×3	97.6	0.45	99.9	0.07	100.0	0.00
C 2×2	91.6	0.58	100.0	0.02	--	--
C 3×2	74.4	0.90	91.2	0.50	100.0	0.04
P 2×3	86.5	2.89	100.0	0.00	--	--
P 3×2	69.3	1.28	80.4	0.86	97.2	0.60
L 3×2	58.7	2.45	72.2	1.94	99.7	0.19
S 2×2	70.4	1.57	100.0	0.00	--	--

Periods 21-30

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	100.0	0.00	--	--	--	--
B 2×3	99.7	0.15	100.0	0.00	--	--
B 3×3	99.6	0.11	100.0	0.02	100.0	0.00
C 2×2	92.2	0.79	100.0	0.00	--	--
C 3×2	89.0	0.61	97.3	0.25	100.0	0.02
P 2×3	89.8	2.46	100.0	0.00	--	--
P 3×2	78.2	0.92	85.4	0.70	98.7	0.38
L 3×2	64.7	2.13	77.3	1.96	99.7	0.21
S 2×2	77.5	1.21	100.0	0.00	--	--

Table 1E: Average transaction price
(by treatment, network, row, and tercile)

Periods 1-10

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	48.0	1.18	--	--	--	--
B 2×3	38.8	1.27	80.5	1.10	--	--
B 3×3	38.9	1.35	81.3	1.84	88.7	2.35
C 2×2	35.7	0.97	83.4	1.09	--	--
C 3×2	18.5	0.73	54.5	1.69	77.8	1.55
P 2×3	90.0	1.62	95.1	0.93	--	--
P 3×2	40.0	1.71	65.0	1.71	79.2	1.27
L 3×2	13.6	0.84	48.6	1.51	66.7	1.89
S 2×2	27.5	1.07	54.8	1.44	--	--

Periods 11-20

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	50.0	0.00	--	--	--	--
B 2×3	49.9	0.06	94.0	0.71	--	--
B 3×3	50.0	0.04	92.9	1.31	95.5	1.83
C 2×2	47.4	0.48	91.2	1.04	--	--
C 3×2	32.8	0.61	75.2	0.86	92.9	0.41
P 2×3	98.3	0.24	98.8	0.21	--	--
P 3×2	60.4	1.28	77.4	1.01	86.6	0.55
L 3×2	19.5	1.51	57.9	2.03	71.2	2.06
S 2×2	37.1	0.98	69.4	1.12	--	--

Periods 21-30

	Row 1		Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
A 1×3	50.0	0.00	--	--	--	--
B 2×3	50.0	0.00	97.3	0.41	--	--
B 3×3	50.0	0.00	97.8	0.63	90.2	2.58
C 2×2	50.0	0.02	92.2	0.64	--	--
C 3×2	46.0	0.48	89.5	0.58	97.6	0.23
P 2×3	99.4	0.09	99.5	0.08	--	--
P 3×2	70.5	0.75	85.2	0.66	91.4	0.42
L 3×2	25.8	2.03	65.4	2.11	76.8	2.05
S 2×2	43.4	0.76	75.9	1.03	--	--

Table 2: Average bids and asks in the asymmetric network
(by trader and period tercile)
# of obs = 80

Node	Periods 1-10		Periods 11-20		Periods 21-30	
	Ask/Bid	Mean	Ask/Bid	Mean	Ask/Bid	Mean
11	Bid	43.58	Bid	65.88	Bid	82.36
	Ask	52.78	Ask	63.59	Ask	68.05
12	Bid	44.80	Bid	63.64	Bid	76.89
	Ask	52.91	Ask	59.30	Ask	66.46
21	Bid	54.41	Bid	73.76	Bid	84.51
	Ask	87.85	Ask	99.10	Ask	99.98
22	Bid	54.21	Bid	72.33	Bid	80.64
	Ask	86.05	Ask	99.21	Ask	99.94

Table 3: Average bids, asks and transaction prices in the asymmetric network
(by path and period tercile)

Tercile	# of obs.	Path	Fraction	Bid	Ask	Price
1-10	54	11 - 21	0.28	62.53	51.80	57.17
		11 - 22	0.22	67.08	48.50	57.79
		12 - 22	0.56	58.74	44.63	51.69
11-20	68	11 - 21	0.34	78.65	58.91	68.78
		11 - 22	0.21	79.29	68.71	74.00
		12 - 22	0.59	74.61	59.94	67.27
21-30	74	11 - 21	0.41	87.33	64.17	75.75
		11 - 22	0.32	84.09	69.36	76.73
		12 - 22	0.35	85.00	65.42	75.21

Table 4: The fraction of completed trades
(by treatment, network, row, and tercile)

Periods 1-10

	Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.
B 2×3	0.9	0.03	--	--
B 3×3	0.9	0.04	0.7	0.06
C 2×2	0.8	0.04	--	--
C 3×2	0.7	0.06	0.7	0.06
P 2×3	0.9	0.04	--	--
P 3×2	0.7	0.06	0.6	0.06
L 3×2	0.8	0.05	0.6	0.07
S 2×2	0.7	0.05	--	--

Periods 11-20

	Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.
B 2×3	1.0	0.01	--	--
B 3×3	1.0	0.02	0.9	0.04
C 2×2	0.9	0.03	--	--
C 3×2	0.9	0.04	0.9	0.04
P 2×3	0.9	0.04	--	--
P 3×2	0.9	0.04	0.8	0.06
L 3×2	0.9	0.05	0.8	0.06
S 2×2	0.9	0.04	--	--

Periods 21-30

	Row 2		Row 3	
Exp.	Mean	Std. Err.	Mean	Std. Err.
B 2×3	1.0	0.02	--	--
B 3×3	1.0	0.00	1.0	0.02
C 2×2	0.9	0.02	--	--
C 3×2	0.9	0.03	0.9	0.04
P 2×3	0.9	0.05	--	--
P 3×2	0.9	0.05	0.8	0.05
L 3×2	0.9	0.05	0.7	0.07
S 2×2	0.9	0.03	--	--

Figure 1: The 3×3 network

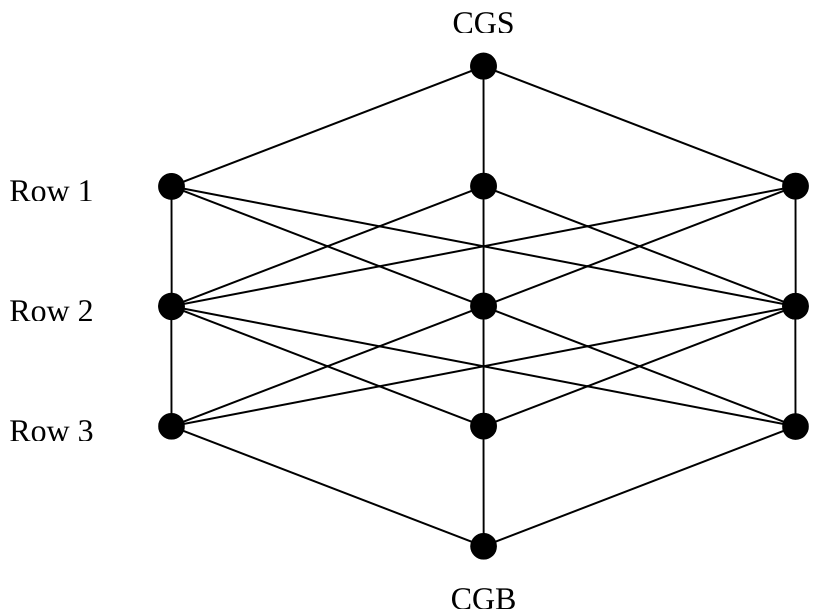

Figure 2: The experimental design

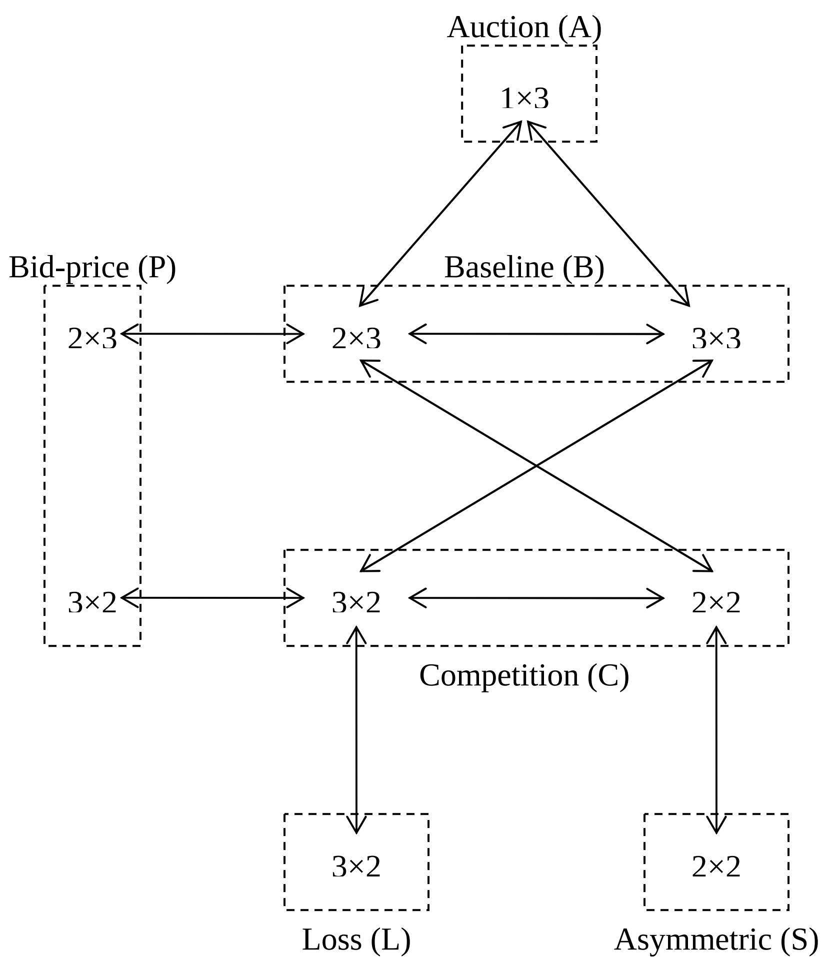

Figure 3A: The effect of intermediation

(Average winning bids in the **A**1×3 network and the bottom rows of the **B**2×3 and **B**3×3 networks, by turn)

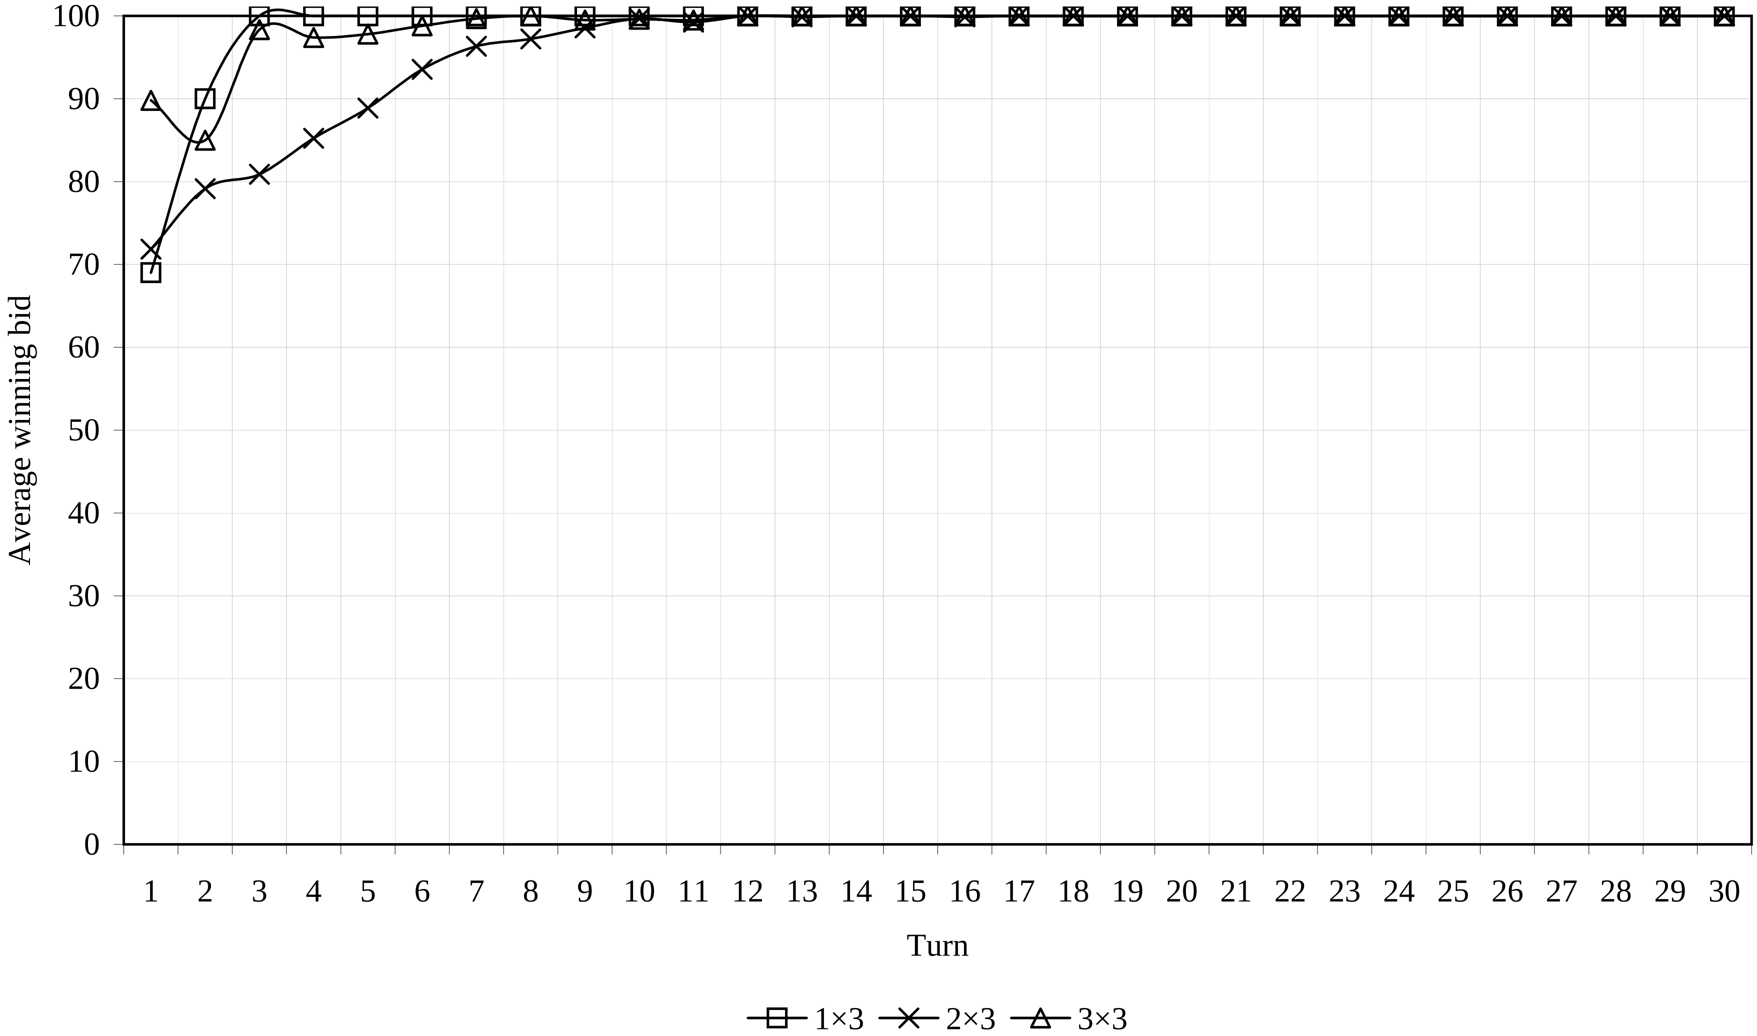

Figure 3B: The effect of intermediation

(Average winning bids in the top row of the **B**2×3 network and middle row of the **B**3×3 network, by turn)

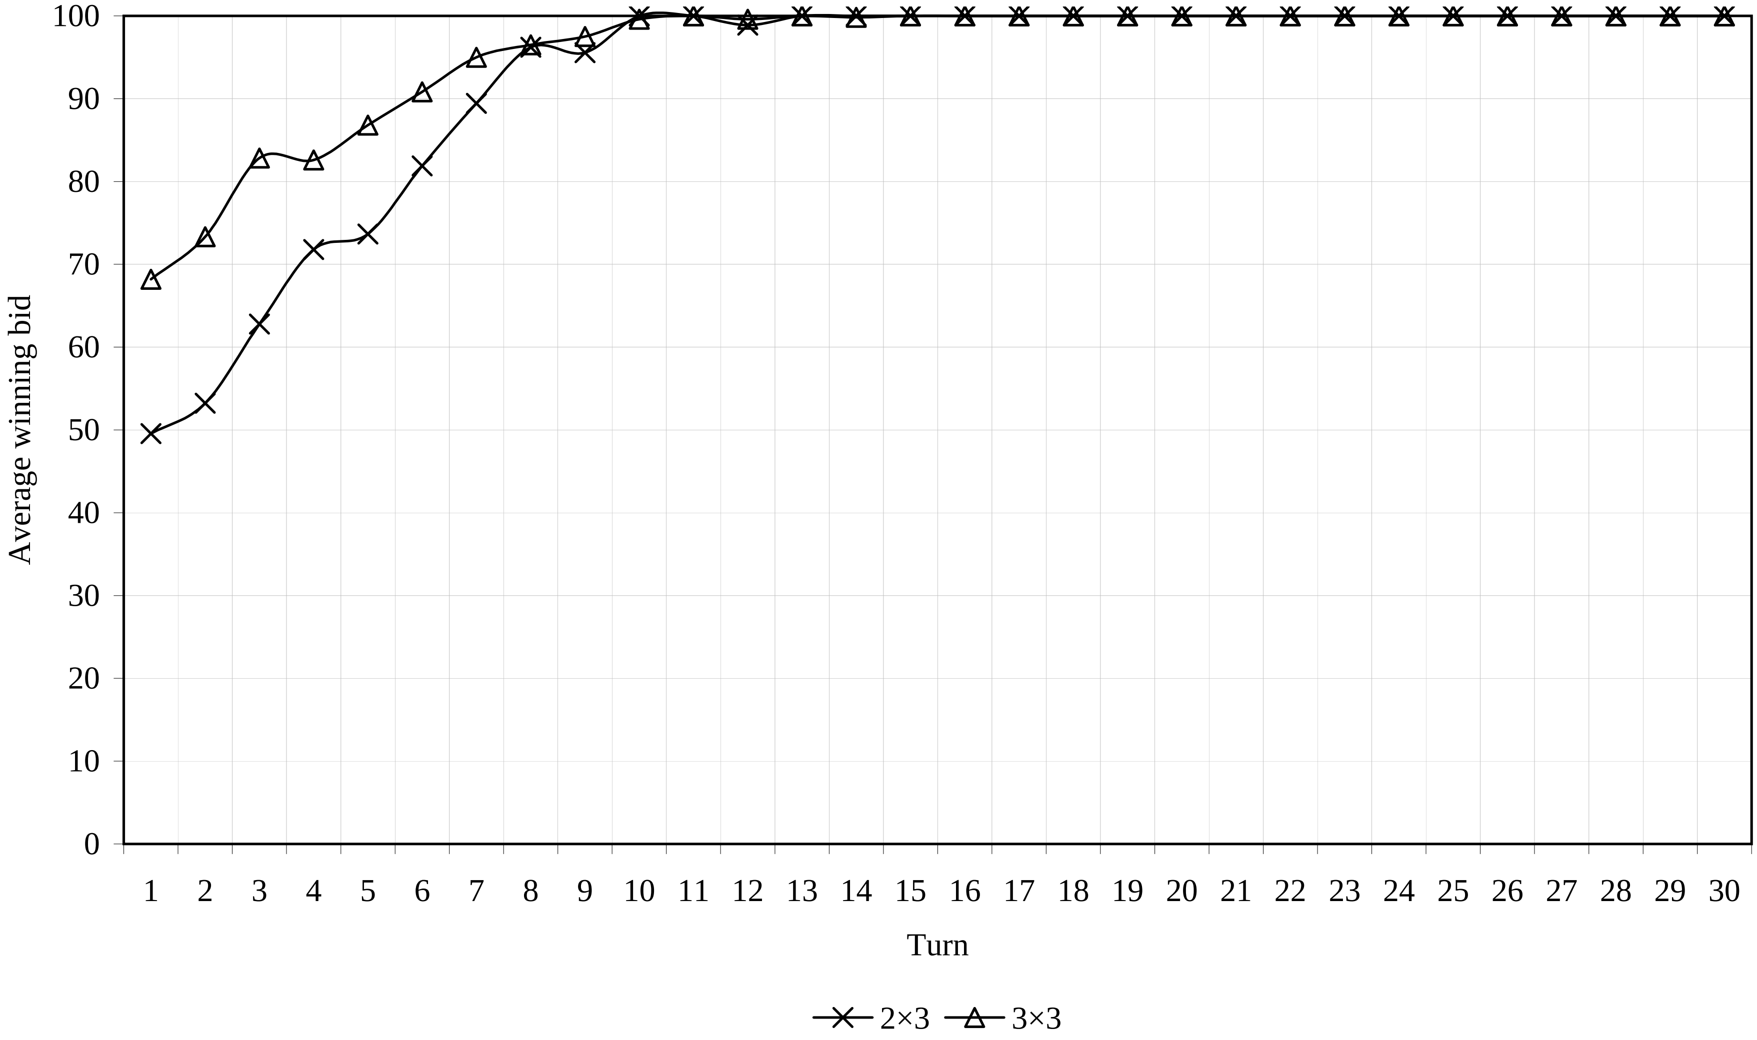

Figure 4A: The effect of competition

(Average winning bids in the top row of the **B**2×3 and **C**2×2 networks, by turn)

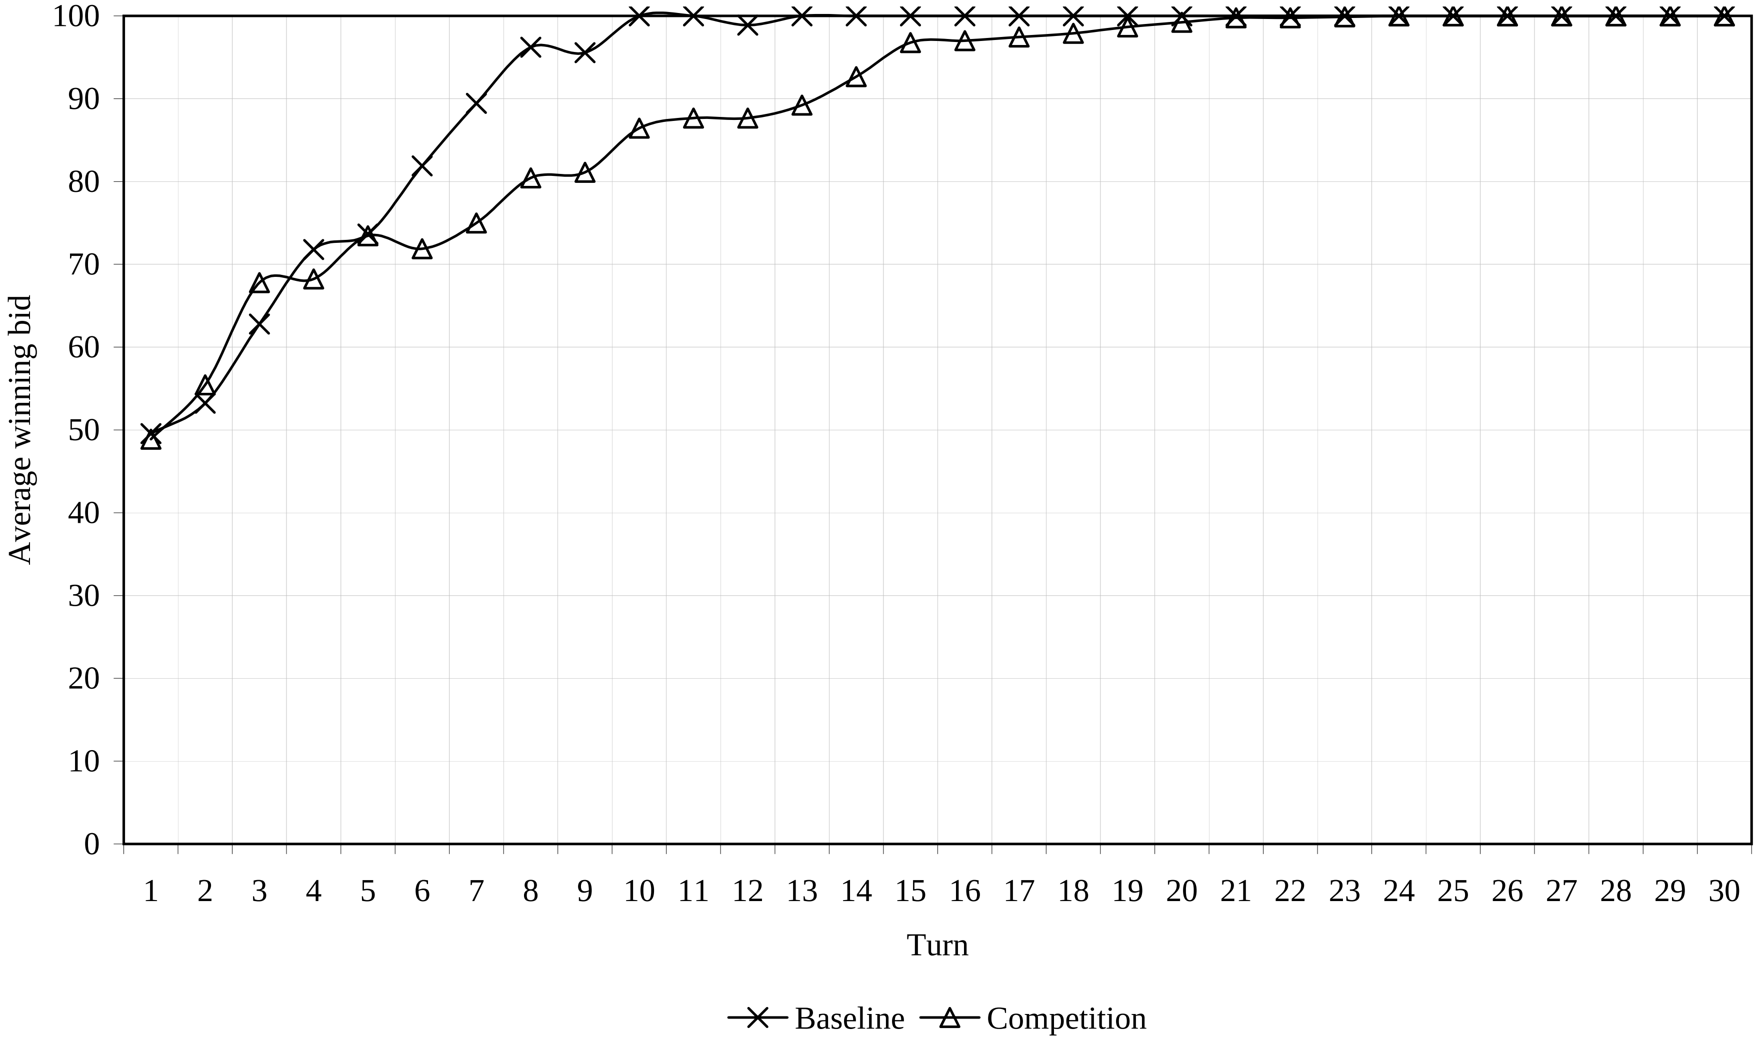

Figure 4B: The effect of competition

(Average winning bids in the bottom row of the **B**2×3 and **C**2×2 networks, by turn)

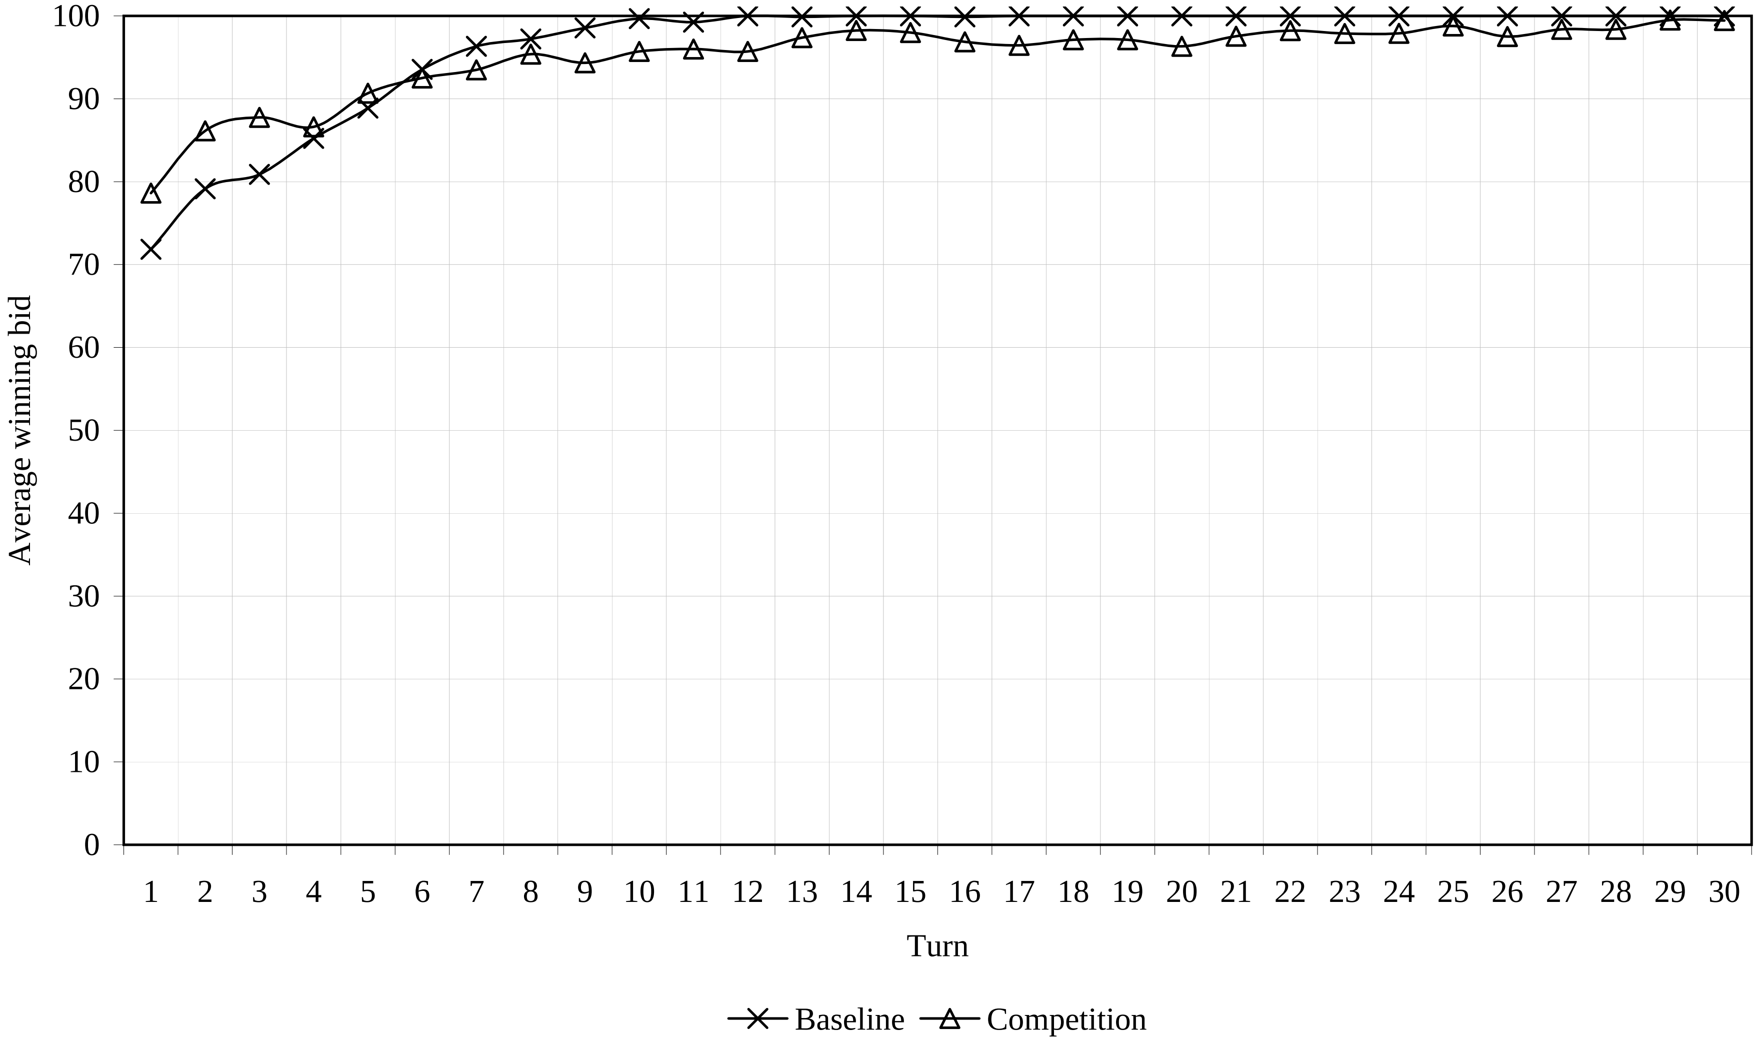

Figure 4C: The effect of competition

(Average winning bids in the top row of the **B**3×3 and **C**3×2 networks, by turn)

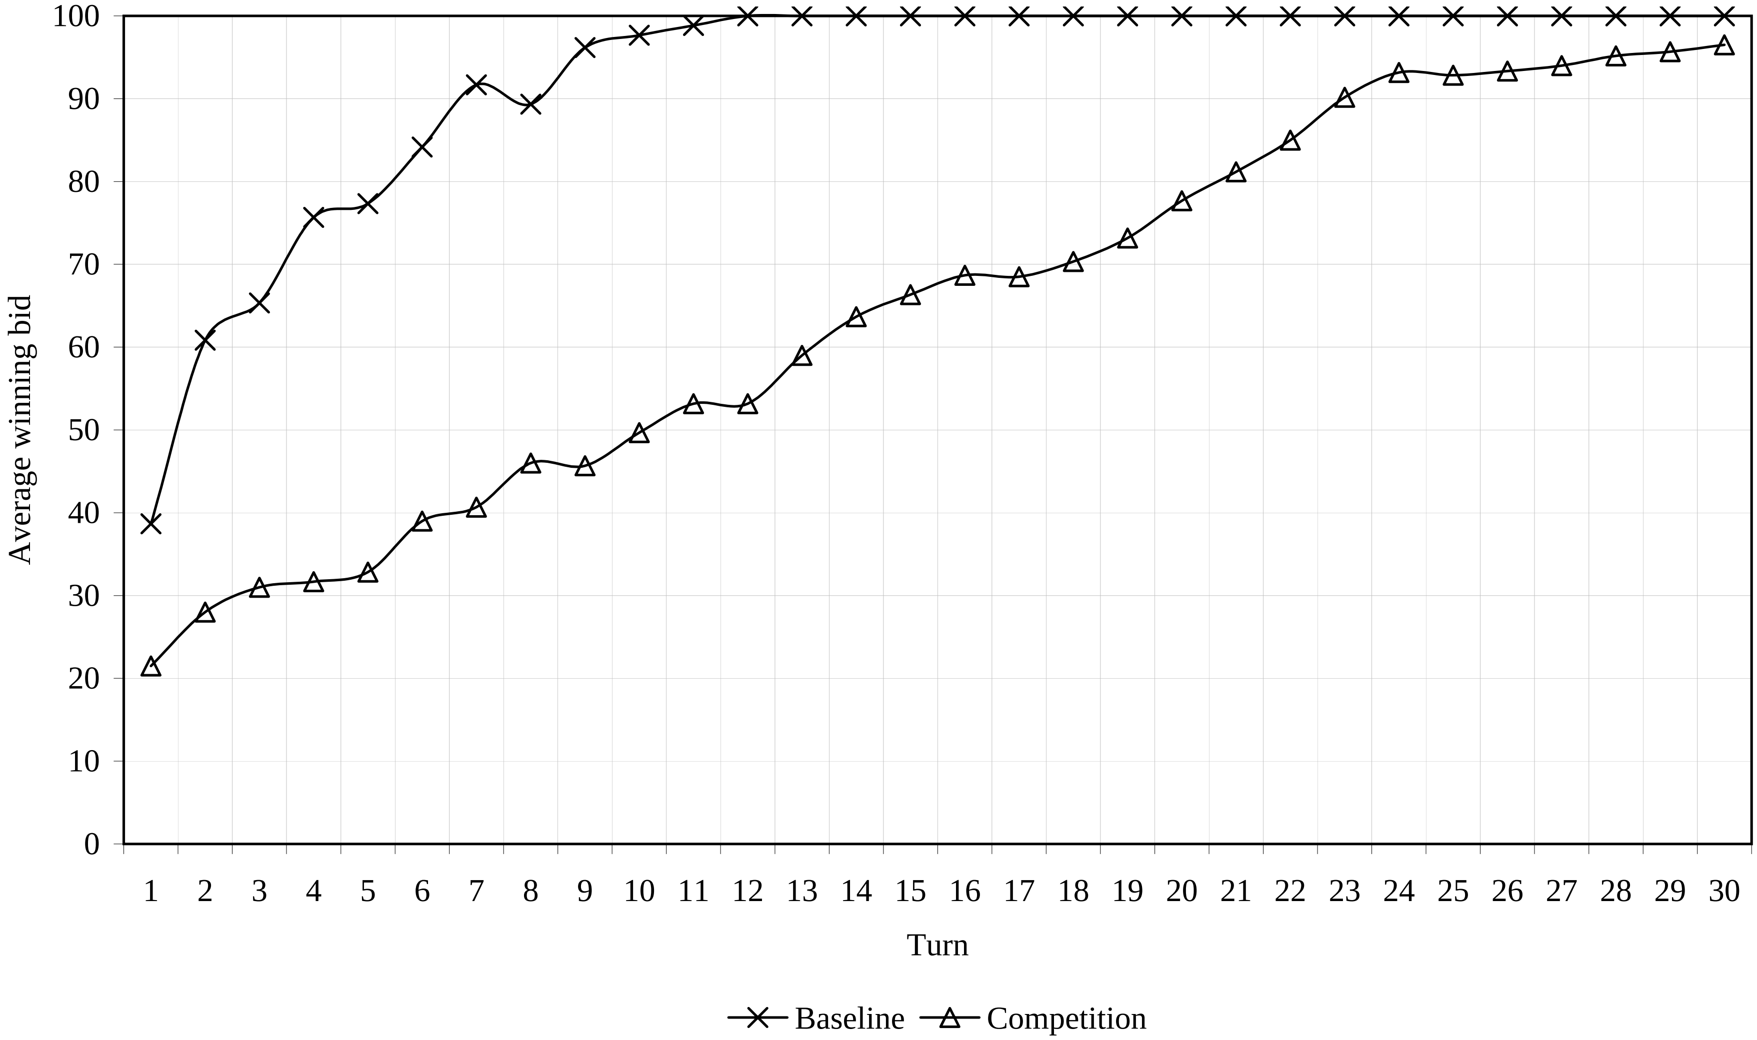

Figure 4D: The effect of competition

(Average winning bids in the middle row of the **B**3×3 and **C**3×2 networks, by turn)

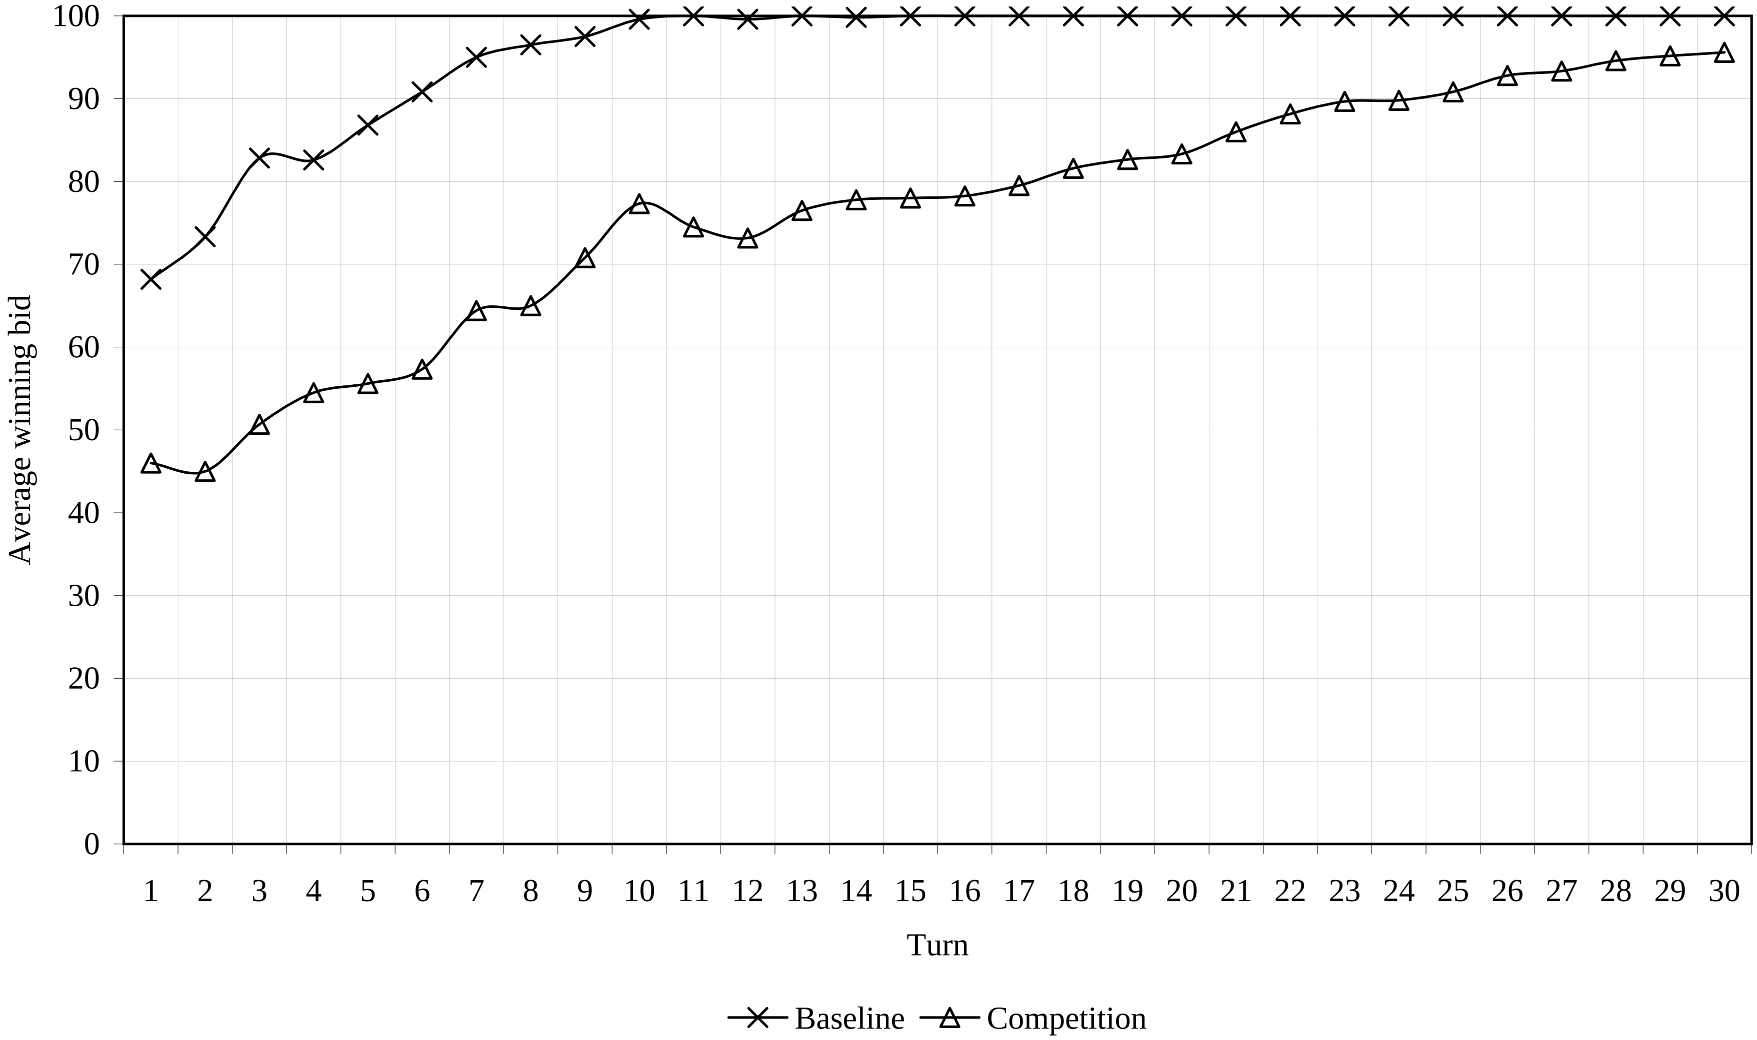

Figure 4E: The effect of competition

(Average winning bids in the bottom row of the **B**3×3 and **C**3×2 networks, by turn)

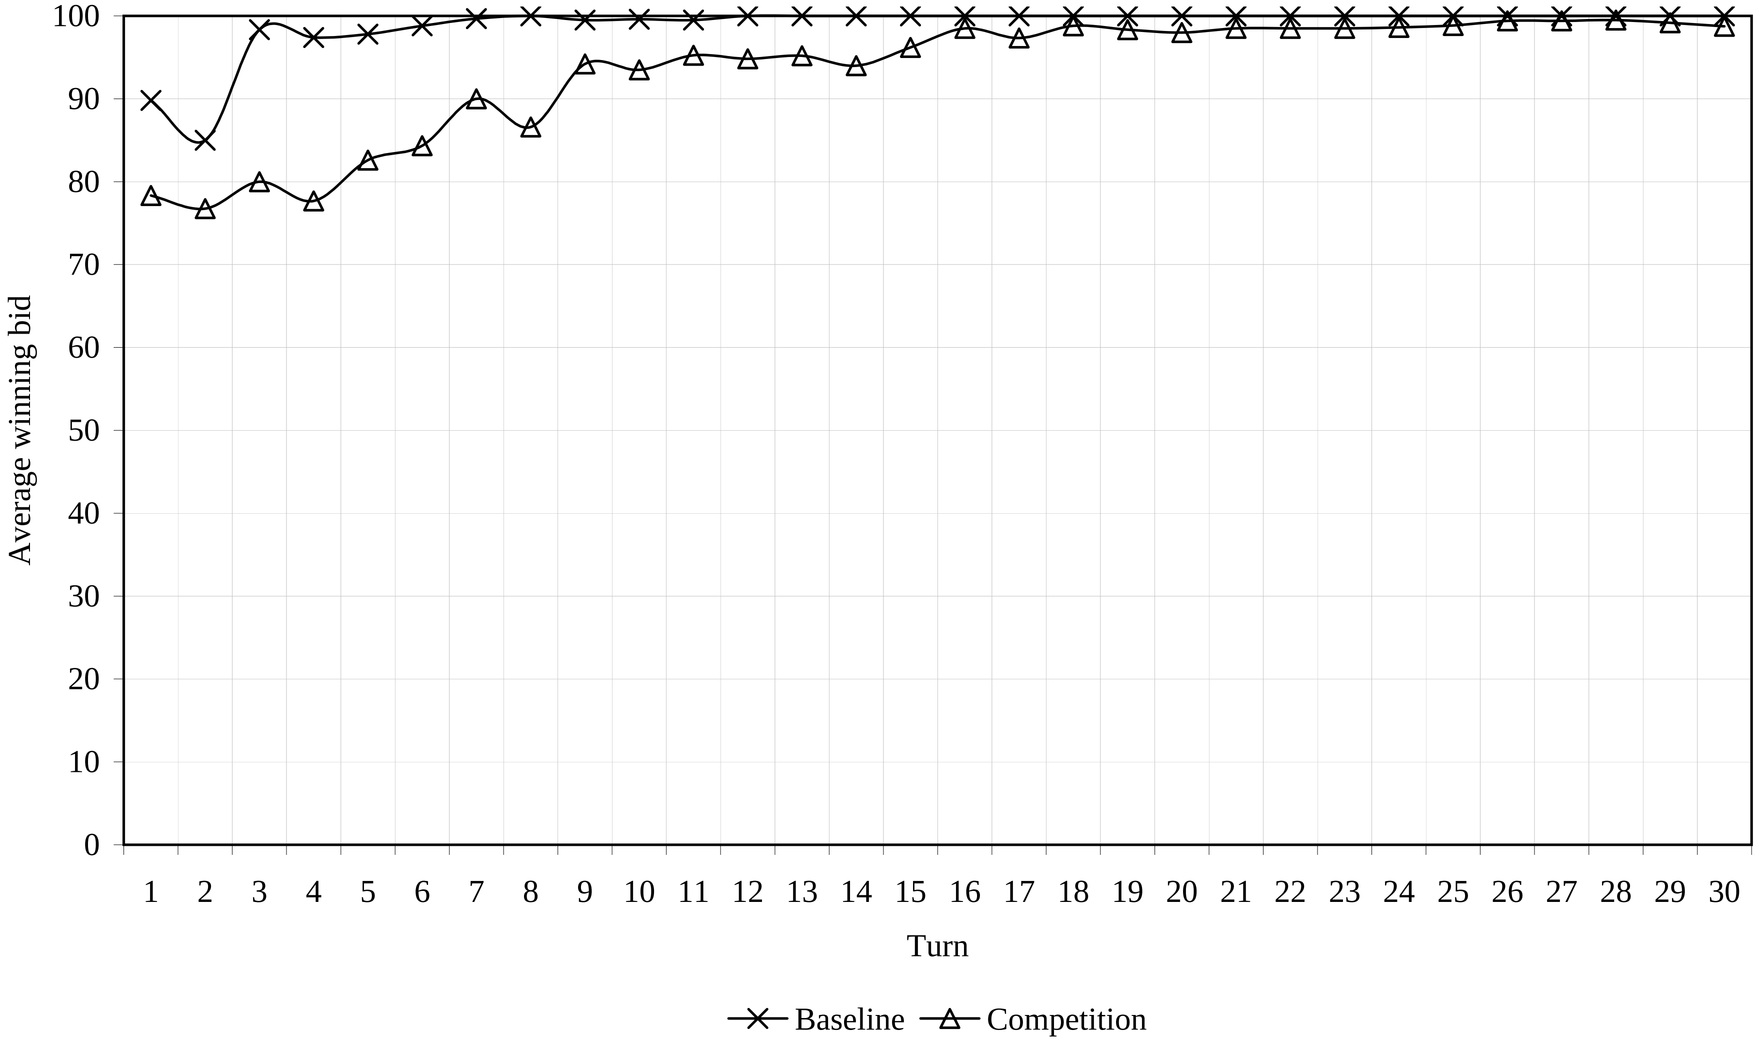

Figure 5A: The effect of pricing rule

(Average winning bids in the top row of the **B**2×3 and **P**2×3 networks, by turn)

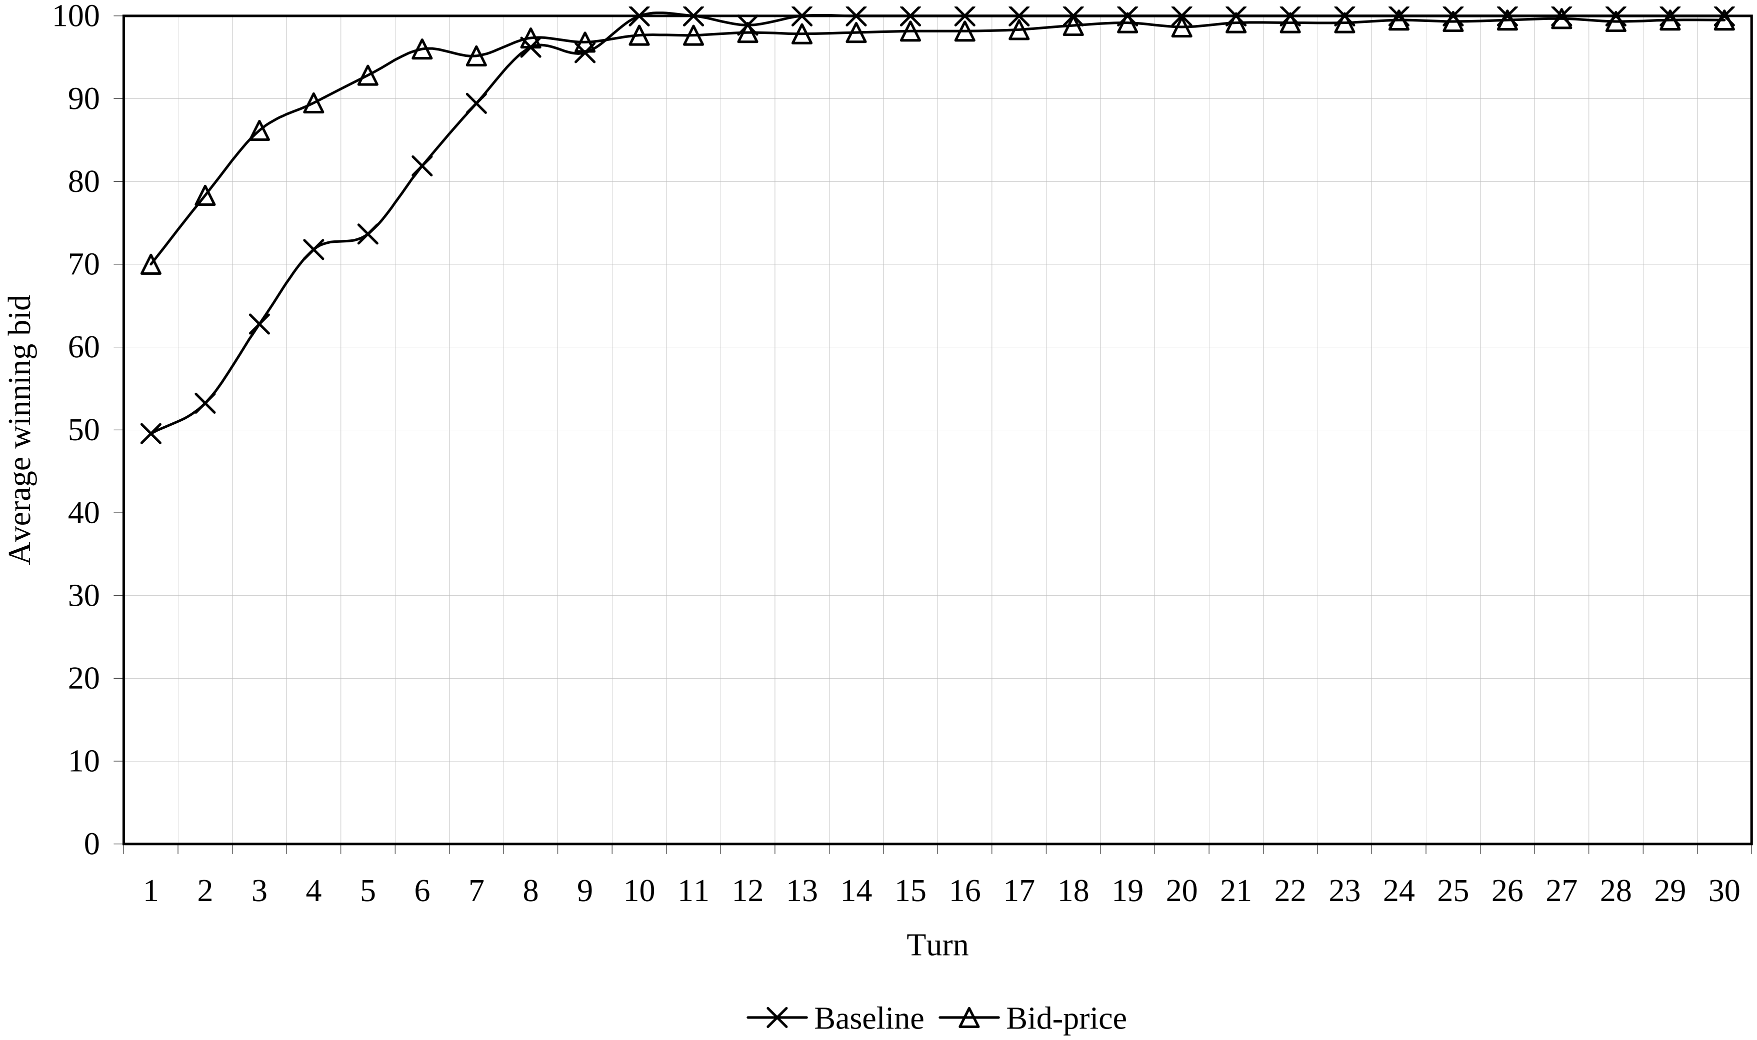

Figure 5B: The effect of pricing rule
(Average winning bids in the bottom row of the **B**2×3 and **P**2×3 networks, by turn)

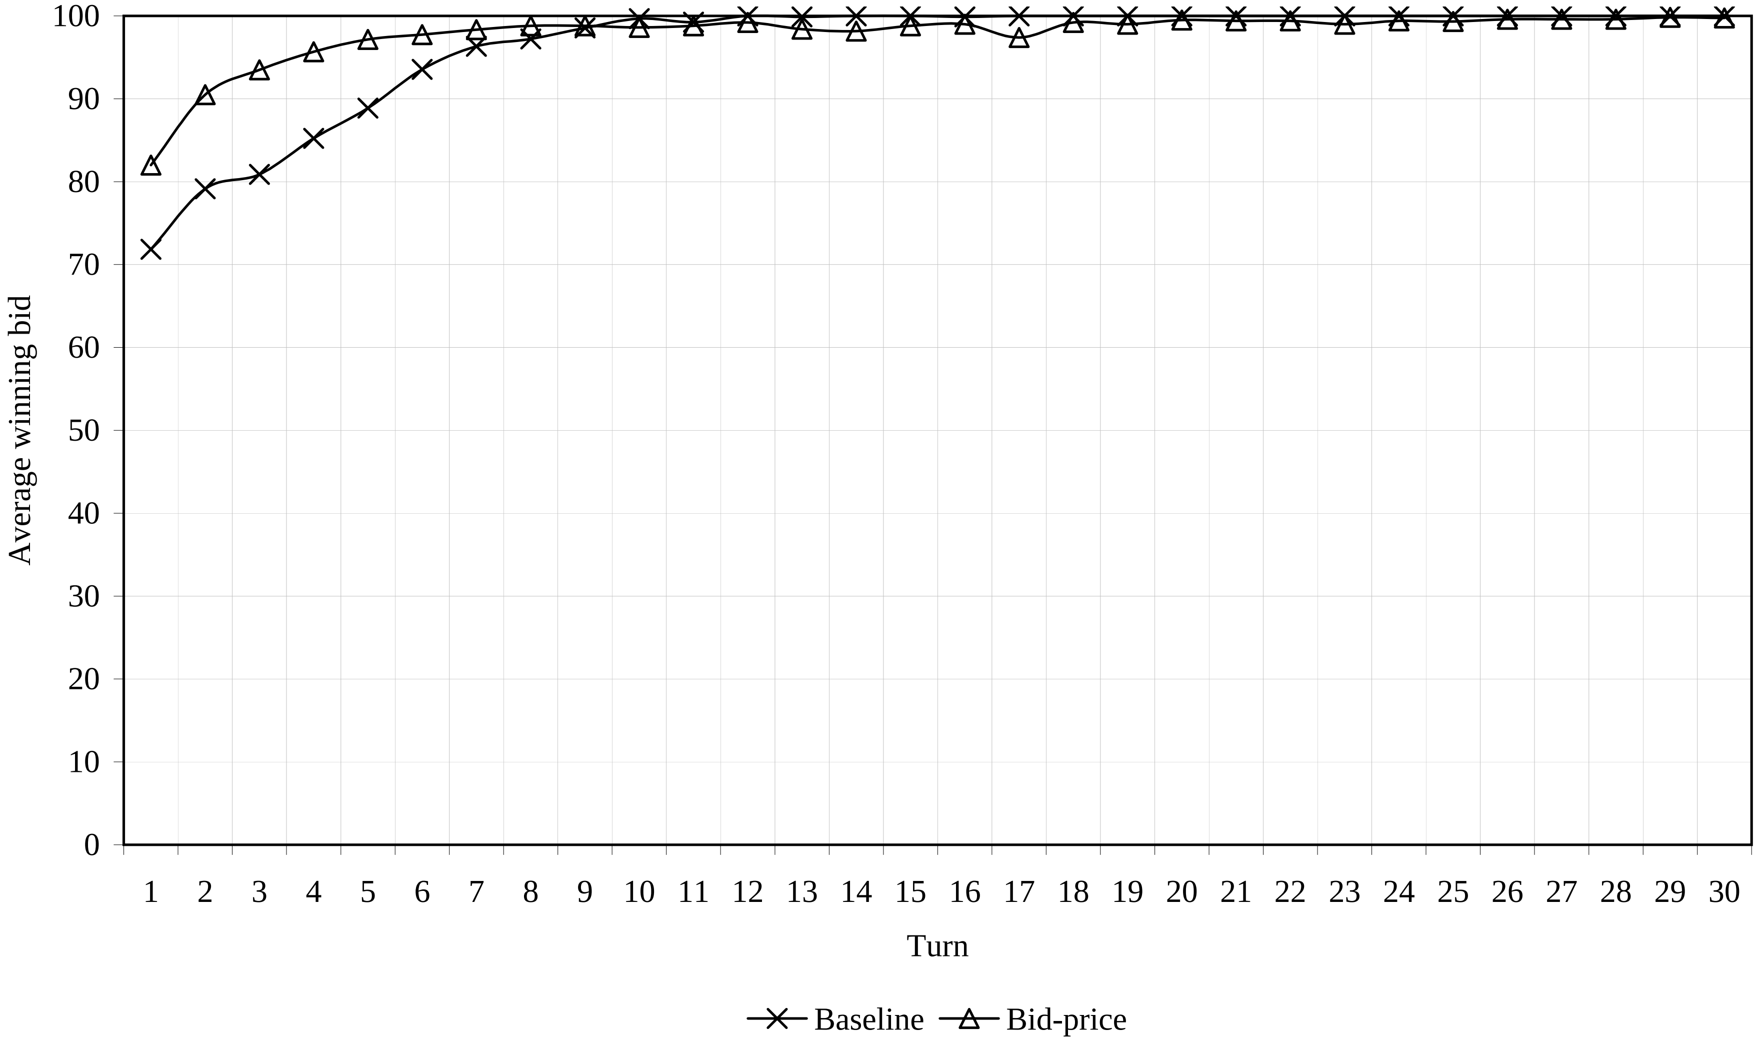

Figure 5C: The effect of pricing rule

(Average winning bids in the top row of the **C**3×2 and **P**3×2 networks, by turn)

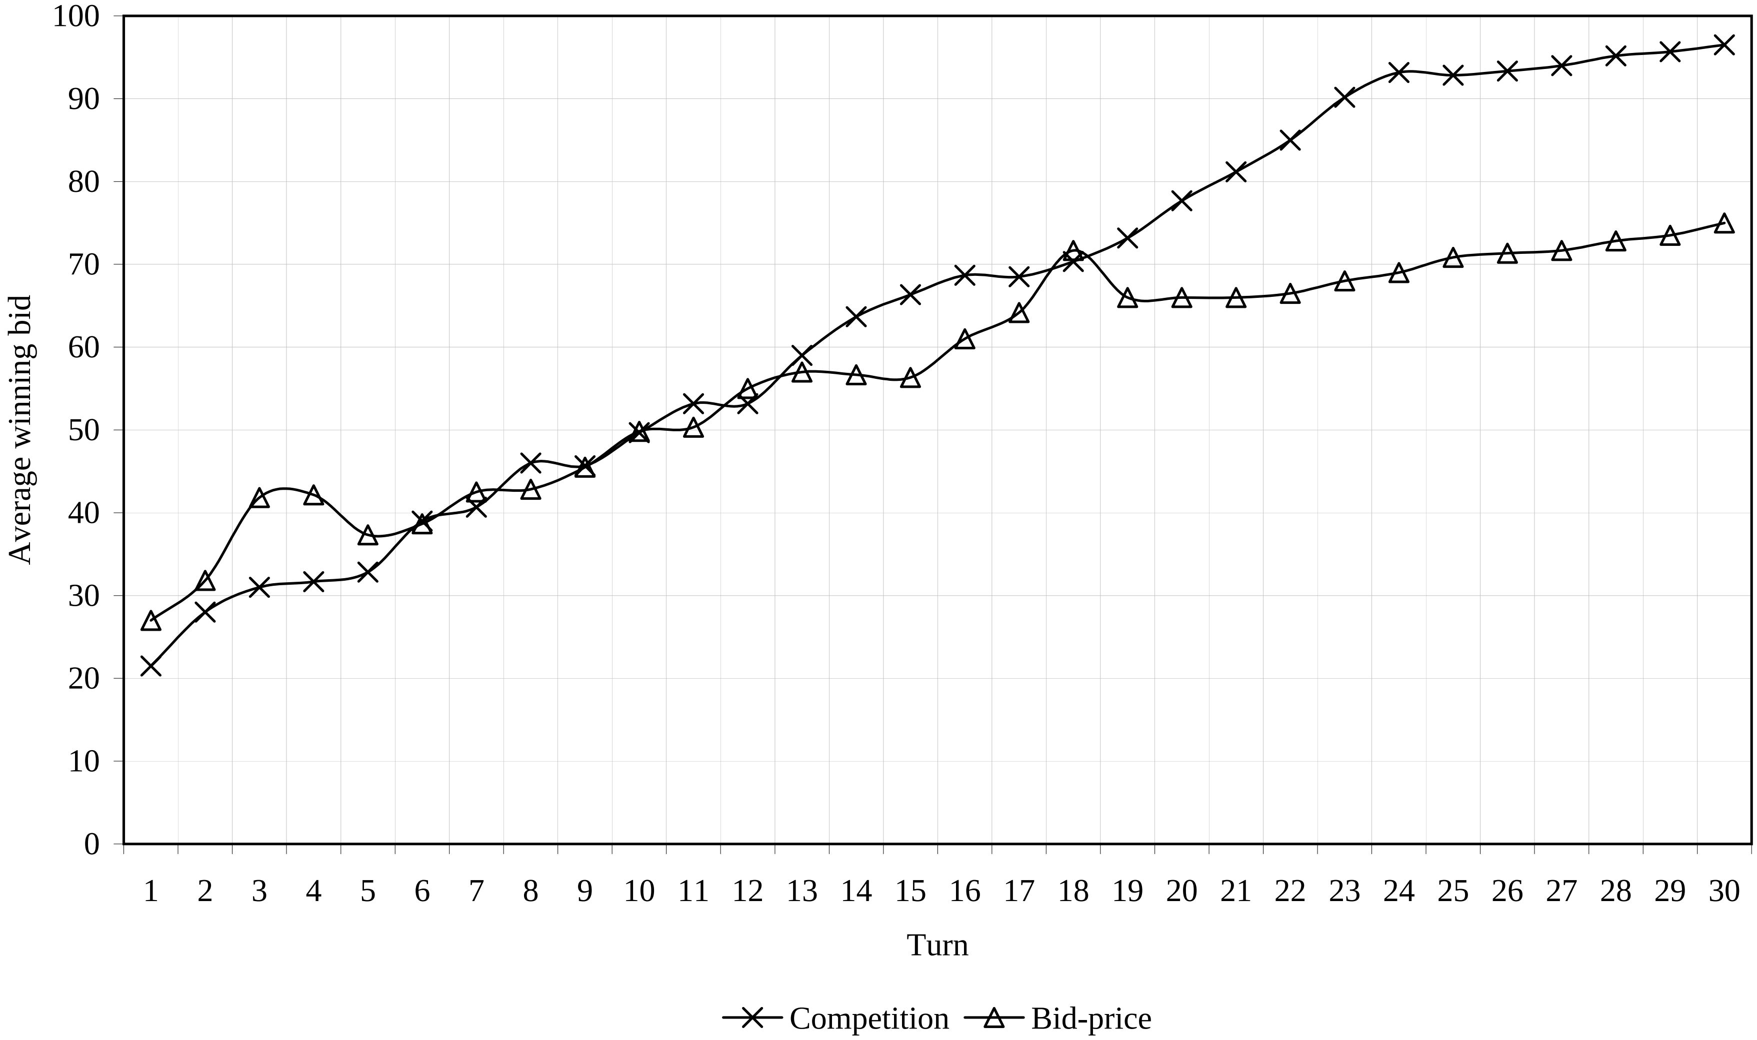

Figure 5D: The effect of pricing rule

(Average winning bids in the middle row of the **C**3×2 and **P**3×2 networks, by turn)

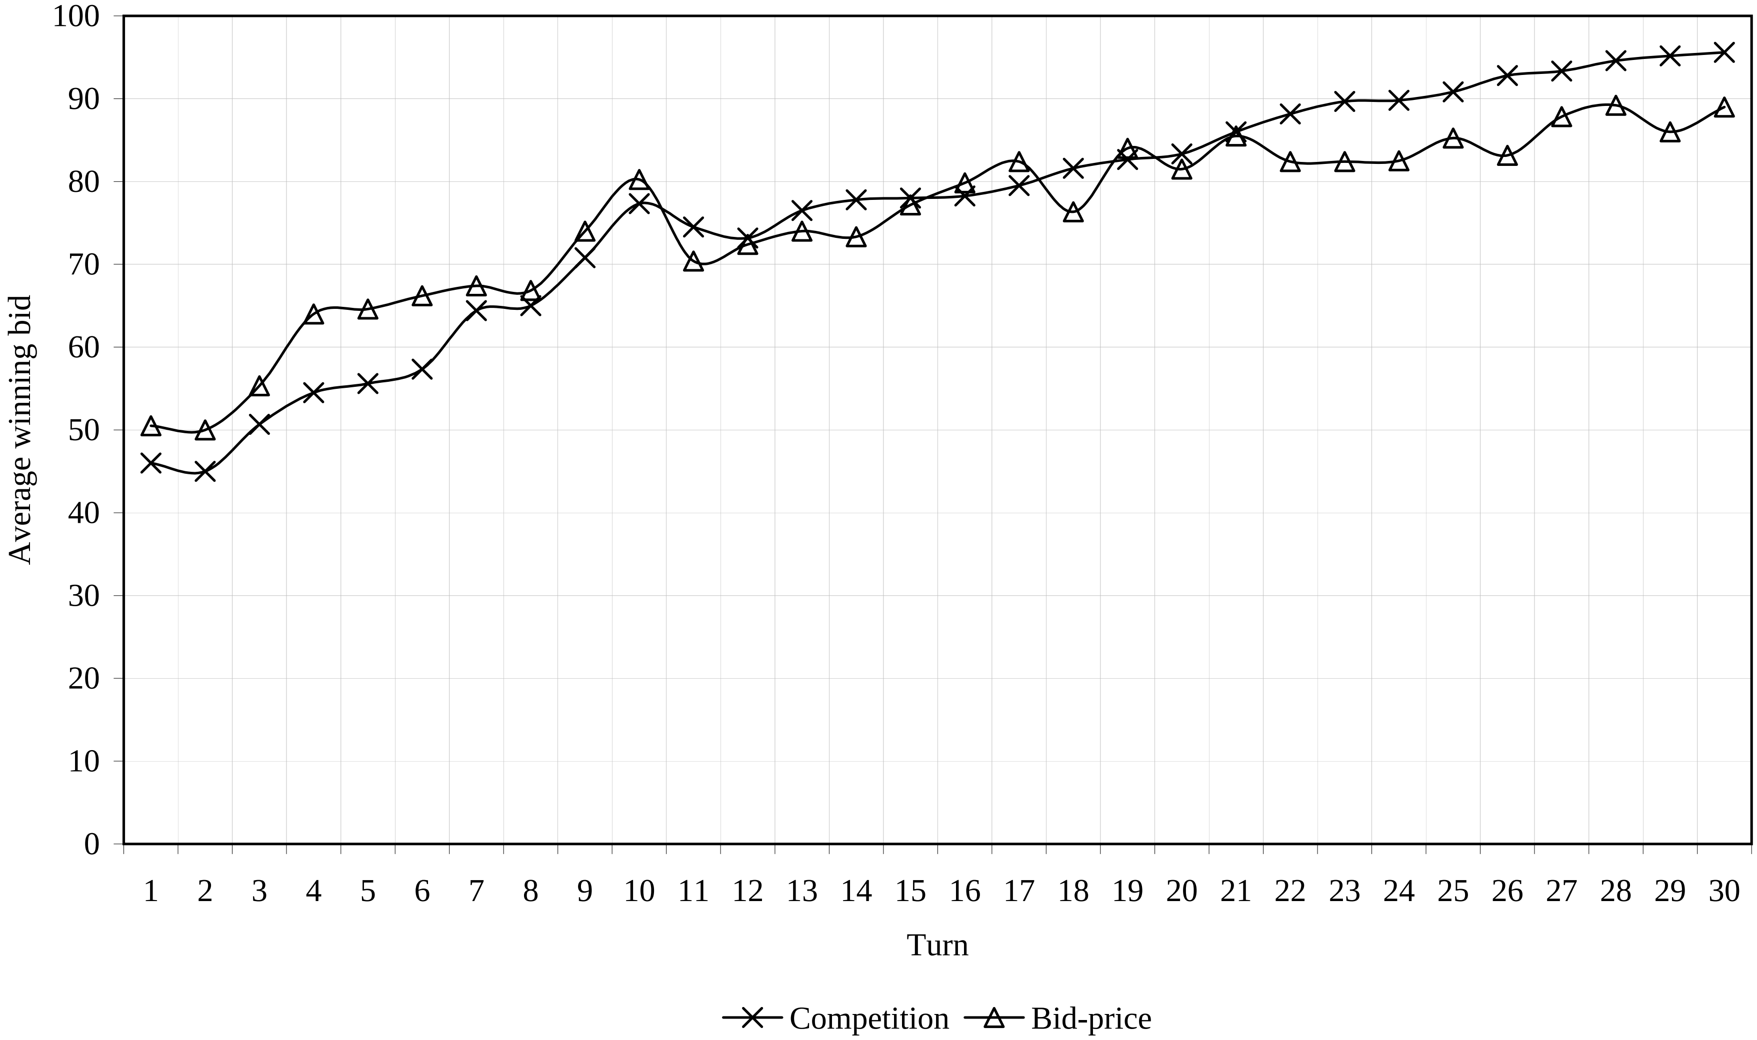

Figure 5F: The effect of pricing rule

(Average winning bids in the bottom row of the **C**3×2 and **P**3×2 networks, by turn)

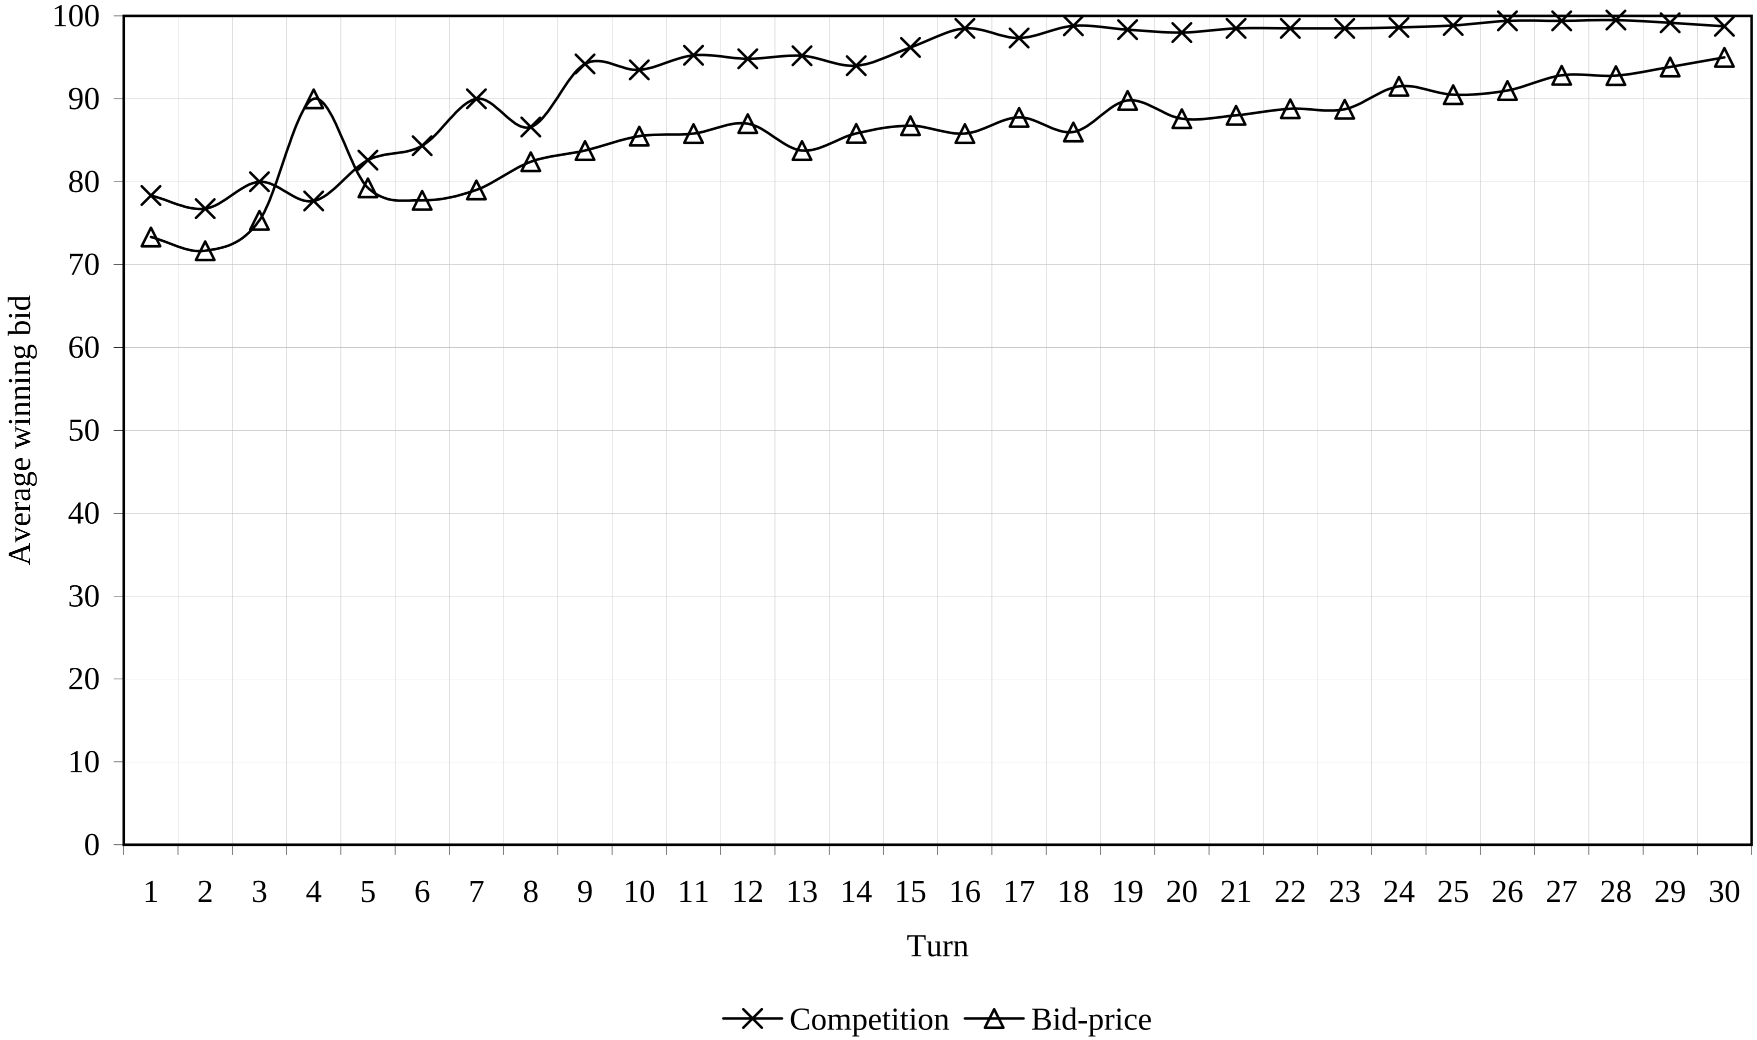

Figure 6A: The effect of losses

(Average winning bids in the top row of the **C**3×2 and **L**3×2 networks, by turn)

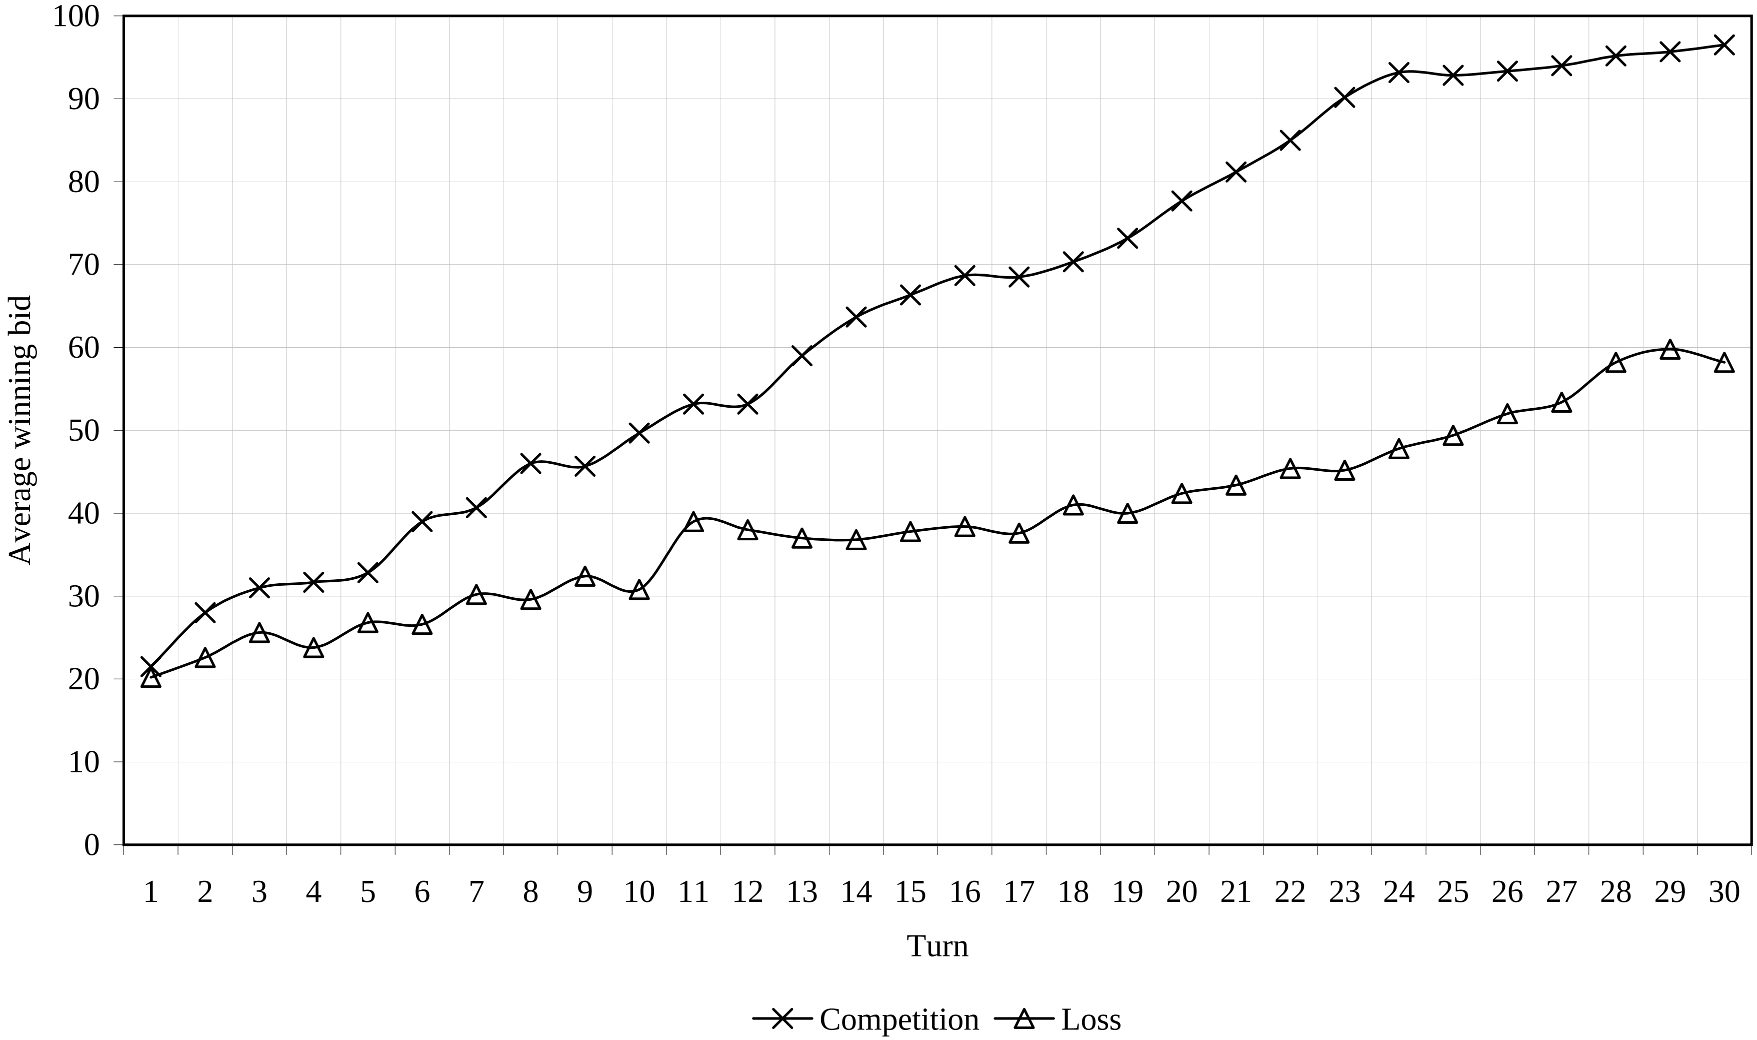

Figure 6B: The effect of losses

(Average winning bids in the middle row of the **C**3×2 and **L**3×2 networks, by turn)

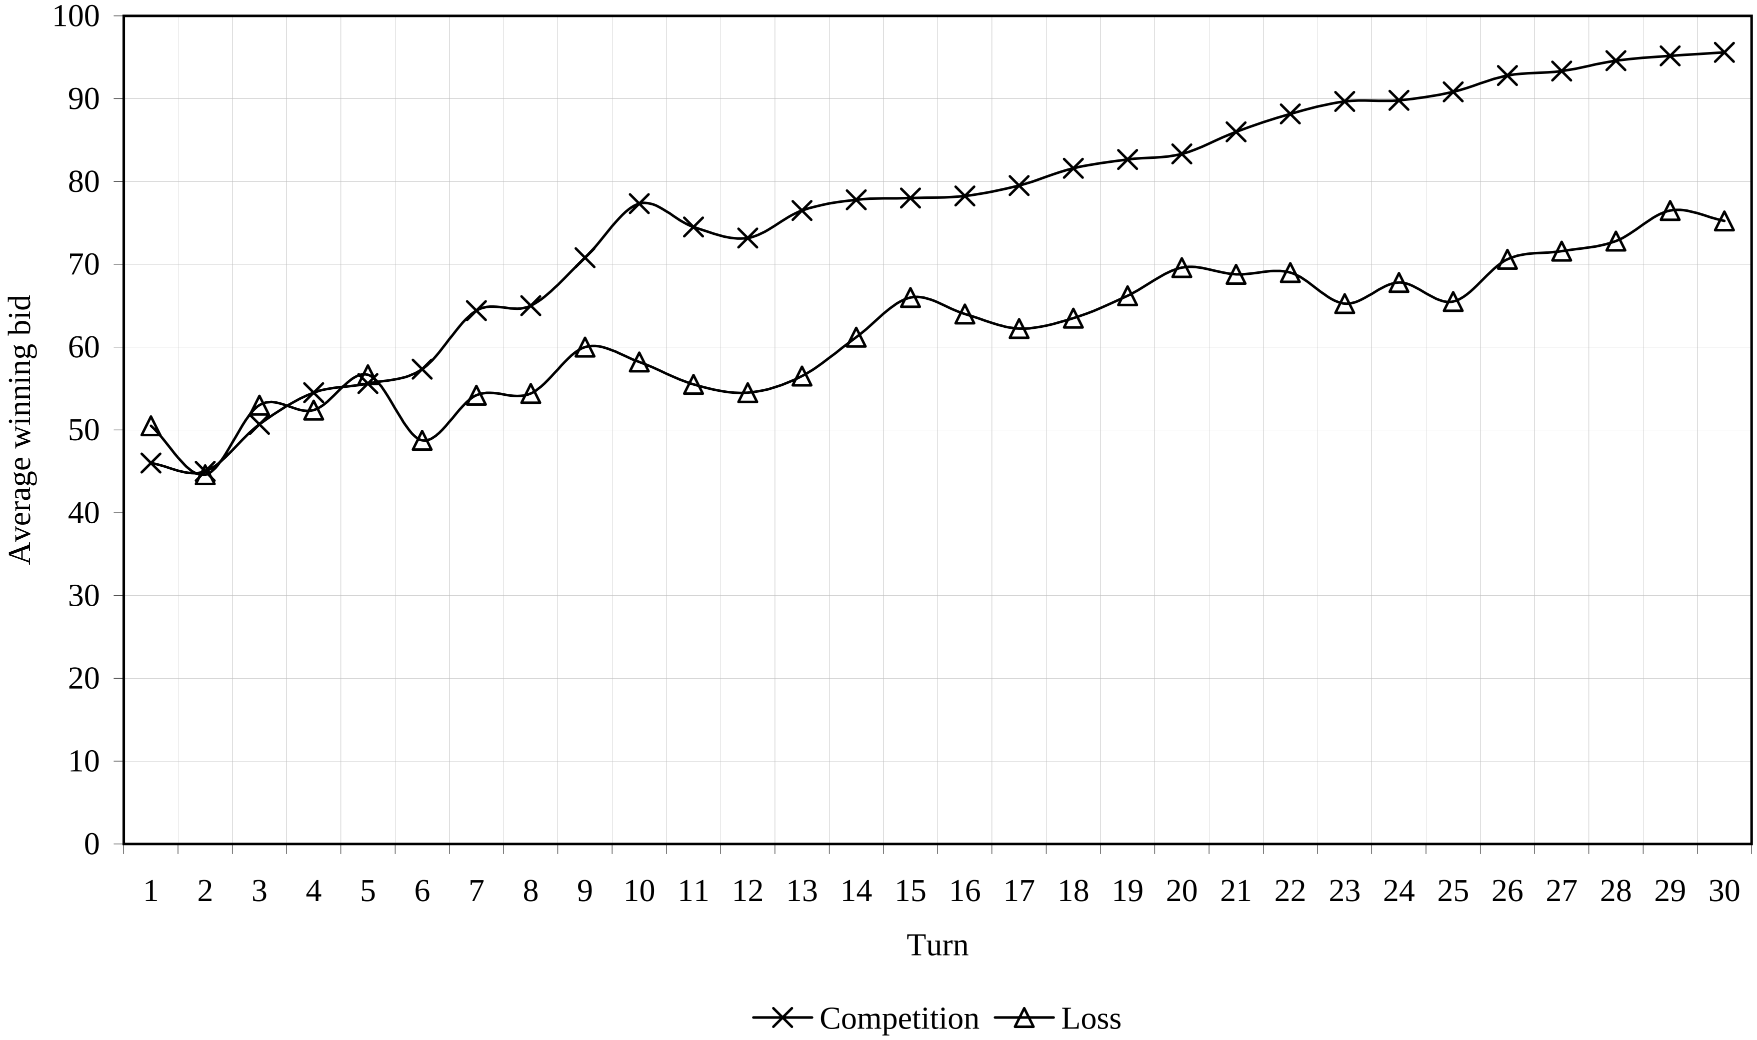

Figure 6C: The effect of losses

(Average winning bids in the bottom row of the **C**3×2 and **L**3×2 networks, by turn)

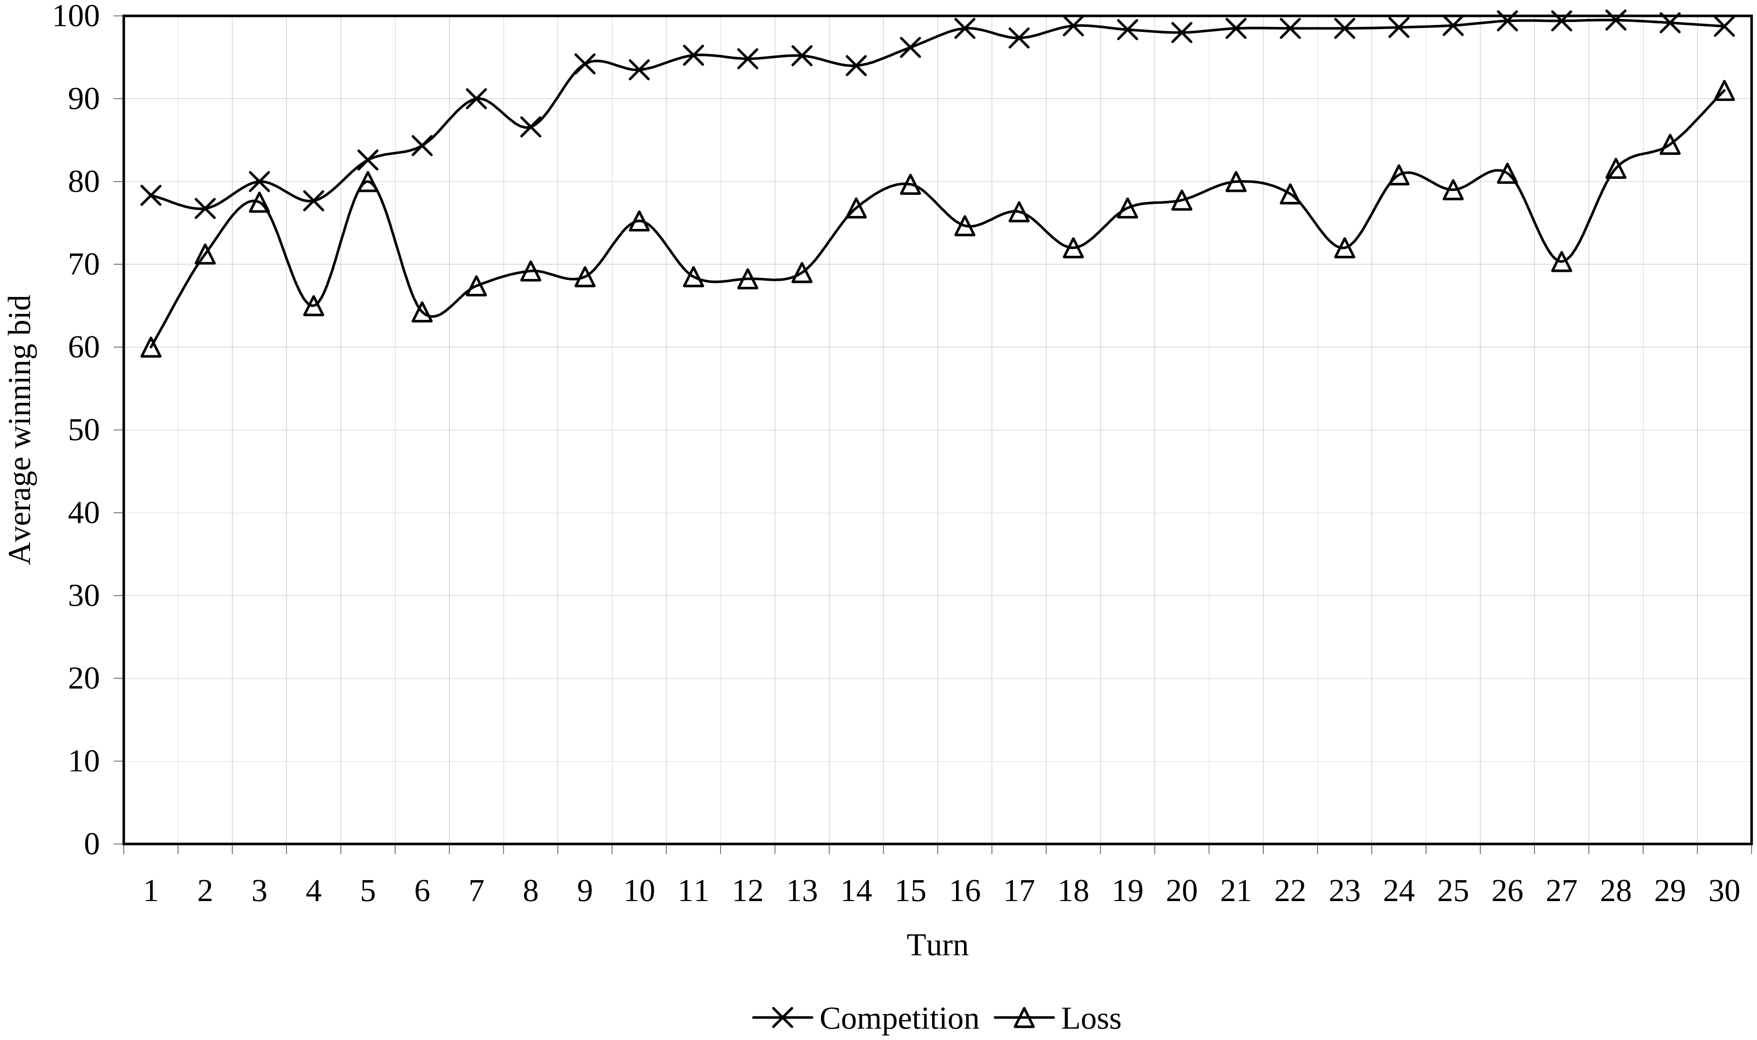

Figure 7: The symmetric (right) and asymmetric (left) 2×2 networks

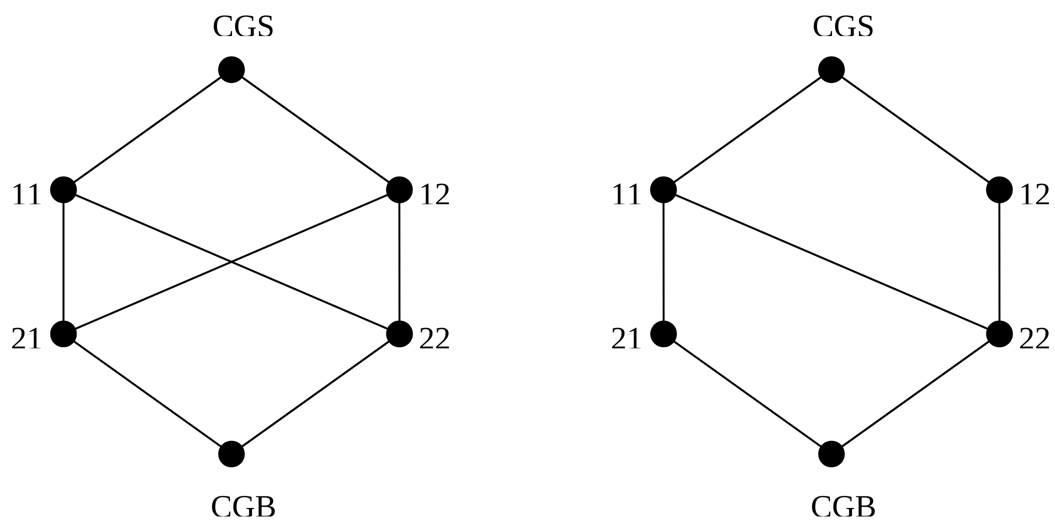

Symmetric

Asymmetric

Attachment I

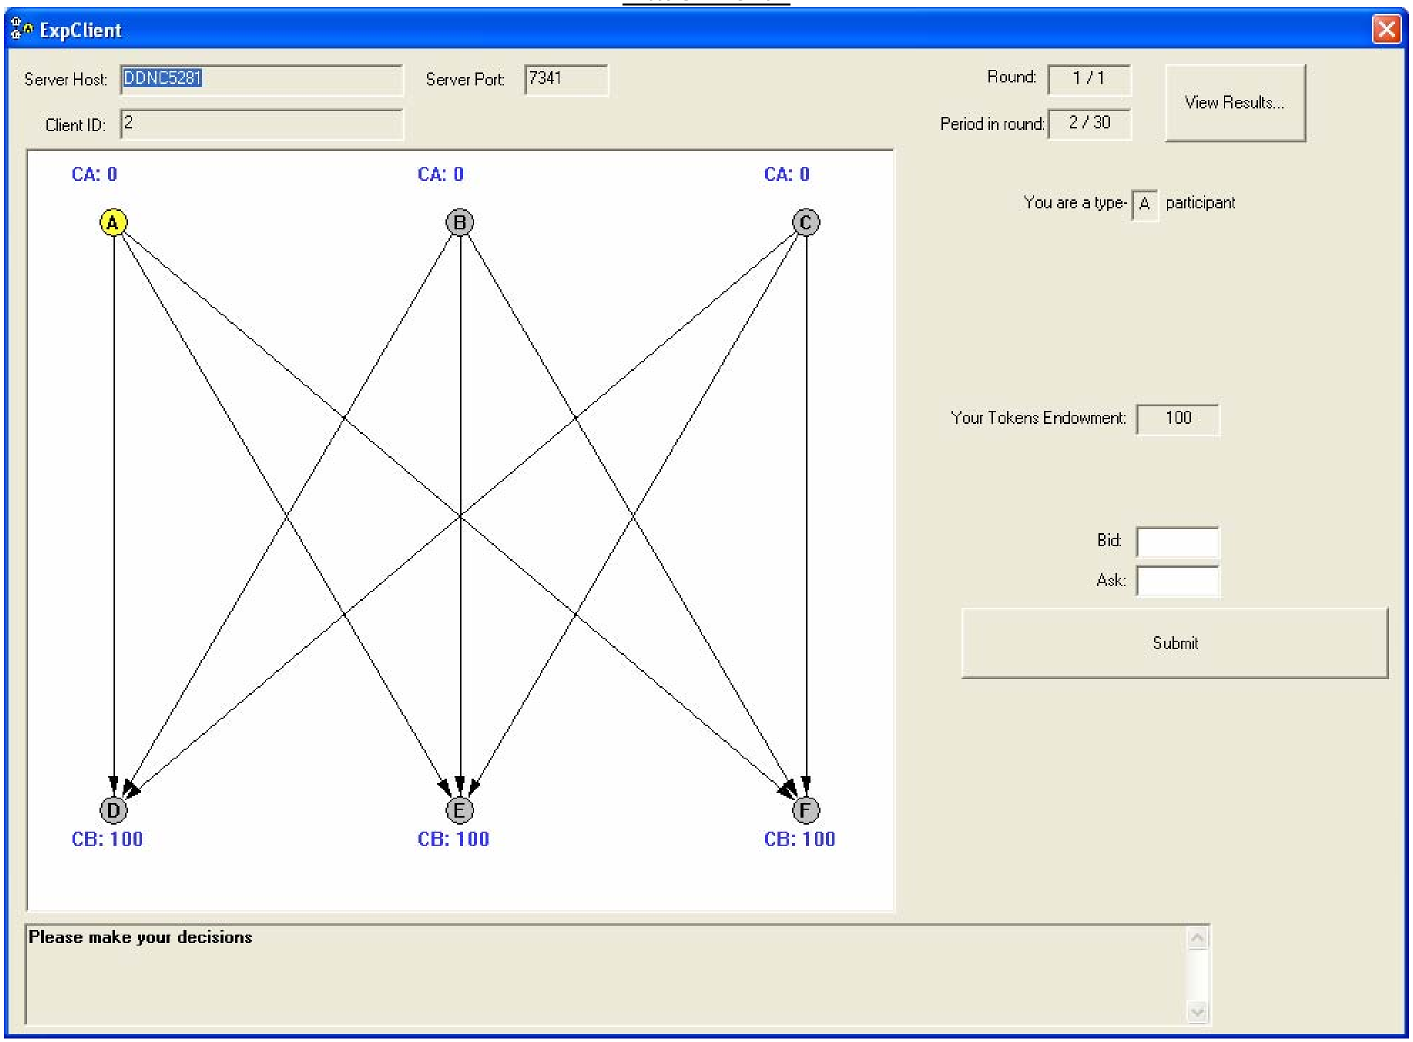

Attachment II