Game Theory and Real Estate Practice in Final Offer Arbitration

Abstract

Final Offer Arbitration, commonly employed in real estate contracts, has been widely studied but primarily in the area of labor disputes. This paper applies the game theoretical models developed in the literature, particularly the work of Brams and Merrill, to the area of real estate, and reviews the results of an experiment in rent determination conducted with real estate professionals. The results of the experiment show consistently sub-optimal behavior by the bidders, which the paper attempts to explicate through ideas developed in the areas of utility theory and mean-variance analysis. The author provides some conclusions designed to assist practitioners.

Introduction

Final Offer Arbitration (FOA) is a technique widely used to resolve disputes in business, and is particularly prevalent in disagreements in real estate. In a survey of 103 real estate brokers at the author's firm 42% had engaged in negotiations governed by FOA, and 16% had done such negotiations five or more times. In FOA (also known as "baseball" arbitration) each party submits an offer to an arbitrator, who must choose the bid that is closer to the arbitrator's perceived "correct" solution.

Oddly enough, the same survey revealed a distinct dichotomy on the reasons that baseball arbitration was regarded as being superior to so-called Conventional Arbitration (CA) in which an arbitrator is allowed complete freedom in selecting an outcome. When asked to choose a reason that FOA was superior (among those who considered FOA better), 64% chose this statement:

> Baseball arbitration is better than conventional arbitration because it makes the parties put forward less extreme positions than conventional arbitration, and therefore causes the parties to converge to a solution.

On the other hand, 26% selected:

> Baseball arbitration is better than conventional arbitration because it threatens both parties, and therefore causes the parties to negotiate a settlement rather than run the risk of submitting to the arbitration.

In discussions with professionals that employed these methodologies, it was surprising to see that there was widespread disagreement about the effect of this common procedure on the bidders and arbitrators who employ it.

There has been broad research and publication on the subject of arbitration methodology, with one thread consisting of the application of the principles of game theory to procedures such as CA and FOA. In the pioneering work of Farber (1980) and Brams and Merrill (1983) the authors developed a game-theoretic model of the optimal behavior of two disputants and a neutral arbitrator.

Brams and Merrill derived optimal positions for FOA based upon the assumption (which we follow) that the arbitrator's position can be viewed as a normally distributed random variable. Their model's Nash-equilibrium bidding position implies that the players should bid significantly higher (lower) than the hypothesized arbitrator mean position. They conclude in their original article and in their subsequent work, including *Negotiation Game*s (Brams (2003)) that FOA is not, in fact, a convergent methodology.

How do professionals actually confronted with FOA problems behave? The literature on FOA describes empirical data that indicates participants sub-optimize relative to the model derived by Brams and Merrill. If they do sub-optimize, what are the reasons (mathematical, behavioral, or both) for that propensity? And finally, in the field of real estate, do agents—professionals acting as advisors rather than principals—behave differently when acting as advisors to landlords (maximizers) or tenants (minimizers)?

In order to understand these issues better, the author looks at the extension of Brams and Merrill model to include the concept of risk aversion, and the paper illustrates the ways in which ideas from general utility theory can explain behavior like that found in the experiment. As an alternate explanation, the author also attempts to apply ideas from mean-variance analysis to the FOA problems. Both of these approaches suggest that bidding sub-optimally can be explained; the former in the context of lower relative benefit for more extreme outcome and the latter by factoring in the tradeoff between a better outcome and greater uncertainty as to the result.

The former approach suffers from the basic problem of all general utility approaches: it begs the question as to what utility function to employ. Most analysis in general utility focuses on mathematically tractable utility functions, which might or might not have any basis in human behavior.

The mean variance approach, on the other hand, allows the derivation of both and expected outcome of FOA and the variance of that outcome, and the resulting set of mean-variance pairs can be matched to the preference of a "user"—this is particularly useful for an agent who must suggest a bid to a principal and who might not be able to easily derive a utility function from that user. It is far easier to state “for this amount of additional risk, you can expect the following additional benefit” and obtain a decision from a client than it is to attempt to extract a continuous function that describes all of the client’s feelings about incremental return. More importantly, the mean variance approach allows the practitioner to define certain combinations of expected outcome and variance that stochastically dominate others, thereby eliminating whole classes of potential bids.

Structure of this paper

The paper is divided into three parts. It begins with a literature review and a synopsis of the Brams and Merrill model (the BMM). Next we present the results of an experiment conducted to study the behavior of real estate professionals in mock arbitrations. Finally we look at the BMM in the context of utility theory and mean-variance theory, extending its conclusion and relating it to the experimental results.

All of the mathematical formulae in the paper are in Mathematica syntax. Also, we have generally taken the position of the landlord or maximizer in showing results. This was done only because it is slightly more intuitive to most readers to consider profit or utility maximization rather than cost minimization. The results can be generalized easily to include the minimization case.

Literature Review

Research on FOA can be grouped into several areas. First, there is the development of what could be called the *fundamental model* and the solution of that model. This model is presented and developed in two articles that have led to a great deal of additional research. In Farber

(1980), he establishes the game-theoretic framework of the FOA process. He provides a model for the expected utilities of the parties, and solves the model for the Nash equilibrium final offers that the parties should employ to optimize their position. He assumes, as do most others and as we follow in this paper, that the arbitrator can be viewed as an exchangeable random variable with a mean and standard deviation.

Brams and Merrill (1983) extended this work by exploring different distributions for the arbitrator random variable. In this article they found that a wide variety of distributions could be employed and solved. In particular unimodal and symmetric distributions (such as the normal) were thoroughly analyzed.

A second area within this research concerns the topic of *risk aversion*, which was incorporated into Farber's model in the form of a utility function. In the interests of defining necessary and sufficient conditions for various distributions, this utility function was not in the Brams and Merrill model.

Farber observes that a risk-averse party will make less extreme offers and will win in the arbitration more often. These results were empirically and mathematically studied by Ashenfelter and Bloom (1984) who examined two sets of arbitrations stemming from NJ municipal labor negotiations, one set settled via CA and the other settled via FOA. In the FOA set, they observe that union offers are accepted more frequently than employers, when symmetrical behavior by the protagonists and fair judgment by the arbitrator would suggest an equal distribution of outcome between the two parties.

They evaluated and rejected the idea that the arbitrators were biased, instead concluding that the union side consistently sub-optimized their bids and therefore prevailed more often. Various explanations for this have been put forward by Ashenfelter and Bloom (1984) and by Brams (2003).

The sub-optimizing behavior (in the sense that union "victories" in FOA lead to lower wage increases than in CA) could result from a structurally different approach by union officials from the approach adopted by representatives of employers. In Brams and Merrill (1991) they suggest a "bonus" formulation for winning. This internal bonus can be used to explain the consistently sub-optimal bidding by the unions, which resulted (in the data set under study) in their winning 69% of their FOA arbitrations. The basic concept is that unions derive greater utility from victory (perhaps based upon their constituents' reactions) than do representatives of municipal employers.

Kilgour (1994) makes the point that the zero sum nature of the game fails when the bargainers have utility functions that are not linearly related, for example when one party is fundamentally more risk averse than the other. He goes on to prove that the more risk-averse party will tend to make less extreme offers, and to win more frequently than the more risk-neutral party. He reviews both the data from municipal union negotiations and from Major League Baseball, and finds it paradoxical that management bids more conservatively and wins more often in baseball, but management bids more aggressively and wins less often in union negotiations. Several researchers have commented that the anomalous nature of baseball players (who have overcome enormous odds to get to the major leagues and are paid enormous salaries) might make them effectively risk seekers in their salary negotiations. Faurot and McAllister (1992) hypothesize that more risk averse players might avoid arbitration, while more risk-seeking players go to the end of the process. He is somewhat disappointed, however, that a consistent attitude toward risk on the part of players is difficult to reconcile with the empirical data.

Marburger and Scoggins (1996) find evidence that more marginal baseball players are less likely to engage in FOA, and suggest that the MLB results show the self-selection that comes from players with better historical results and therefore less risk of "input substitution." Thus, some players exhibit risk averse behavior and others exhibit risk seeking behavior, based upon their historical performance.

Throughout the literature on negotiation there is a concept of the contract zone—that range of overlapping outcomes that each party finds superior to impasse. Dickinson (2004) shows that the more risk aversion that players have, the broader their contract zones become, and therefore the more likely that they reach a negotiated settlement.

In recent work, the study of risk aversion and bargaining *prior* to FOA have been synthesized into a model by Hanany, Kilgour, and Gerchak (Hanany et al. (2007)). They have generalized the model to include general utility functions and they have solved the Nash equilibrium for a logical bargaining algorithm. They prove several things, most importantly that under a regime of risk aversion (with some simple contraints on the utility function) there must exist an outcome that benefits both parties more than going into FOA. This result is particularly important to the empirical results described in the introduction, because in the author's experience the vast majority of negotiations subject to FOA are settled by prior to the actual arbitration. Hanany et al. (2007) illustrate the solution of the general Nash bargaining problem, in which the goal is to find a solution that maximizes the product of the difference between the expected outcome for each player (presuming that arbitration occurs) and some solution **X**. They establish the existence and uniqueness of this solution, and provide a general method for finding that solution.

This provides a valid framework for understanding sup-optimizing bidding, but presumes (as is the case in most utility theory) that a utility function can be determined either theoretically or empirically.

On a separate track, several researches have attempted to determine whether in practice (either in experiments or in documented cases such as described above) CA and FOA differ markedly in their *convergence*. For example, Ashenfelter et al. (1992) achieved some surprising results when testing CA against FOA. Their data suggested strongly that baseball players are less likely to come to a negotiated solution when FOA is employed than when CA is employed, which is directly contradictory to the conventional wisdom. They also found no strong evidence of risk averse, sub-optimizing behavior under FOA.

Of particular interest to negotiation practitioners is the work of Ashenfelter and Dahl (2005) in which they analyze the data from arbitrations under the New Jersey Fire and Police Arbitration Act over a longer sample horizon than their initial analysis. They find that this longer sample shows a consistent dampening of the asymmetrical result that dominated their 1984 analysis, and they attempt to explain this by analyzing the changing role of *expert agents* in the outcome of bargaining and arbitration. They conclude that expert agents improve the probability of winning for the side that hires them, both by influencing the decisions of arbitrators and by moderating overconfidence on the part of their clients. The increasing prevalence over time of such agents in their sample explains the gradual equilization of outcome between management and labor. However, they also conclude that this advantage over time has transformed into a classic prisoner's dilemma situation, in which parties almost universally employ agents, thereby neutralizing the benefit to both parties.

Brams model

Farber and Brams and Merrill's model both consisted of three steps: a) defining a payoff function for the problem that states the expected outcome to the participants; b) determining a maximum point for that function and c) determining whether that maximum point is locally or globally maximum.

A key to constructing such a payoff function is determining the manner in which the arbitrator is hypothesized to behave (the arbitrator function). While Brams and Merrill studied a wide variety of arbitrator distributions, this paper will focus exclusively on a very tractable and perhaps the most logical distribution, the normal. This requires us to follow the assumption that both parties have similar information about the arbitrator's function, and so we can assume that both parties know the mean and the standard deviation of the arbitrator's position. Real estate, with its well-established databases on comparable transactions, is a good example of a field in which the normality assumption should hold.

First, we state the Brams and Merrill payoff function. Let **a** be the bid of the minimizing party and **b** be the bid of the maximizing party. The payoff (g[a,b] in Figure 1) will be:

$$g[a\_, b\_] := a\,F\left[\frac{a+b}{2}\right] + b\left(1 - F\left[\frac{a+b}{2}\right]\right)$$

Figure 1

Where F is the CDF for the normal distribution, as in Figure 2:

$$F[x\_] := \int_{-\infty}^{x} \frac{1}{\sigma\sqrt{2\pi}}\, e^{-\frac{1}{2}\left(\frac{t-\mu}{\sigma}\right)^{2}}\, dt$$

Figure 2

The μ and σ are the assumed mean and standard deviation of the arbitrator's (normal) preference function.

The payoff function basically states that the expected outcome of the arbitration starts with the probability that the arbitrator's choice is above or below the midpoint of the two bids. The probability that the arbitrator will choose a position below is multiplied by **a**'s bid, and one-minus that probability is multiplied times **b**'s bid. The sum of the two is the expected outcome.

Using some realistic parameters, we can graph the BMM. We assume that FOA is being employed for a real estate rent arbitration. Based upon the analysis of transaction comparables, we assume that the arbitrator's choices will center around a mean rent of $40 with a standard deviation of $10. These mean and standard deviation assumptions will be used throughout the balance of the article.

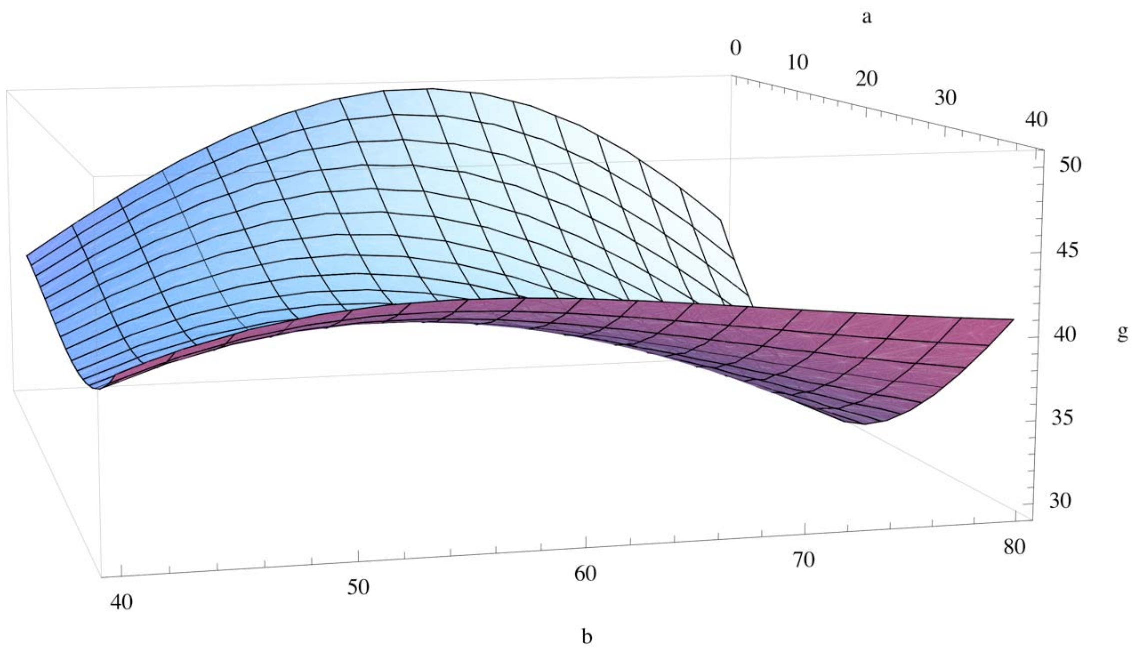

Figure 3: The expected outcome of a FOA arbitration (g) as a function of the bids of the two parties (a and b), illustrating the characteristic saddle shape of a Nash-equilibrium problem

By taking the first derivative of the payoff function **g** and with a little algebraic manipulation, it can be determined that the optimal bids (with **a** being the minimizer and **b** being the maximizer) are as in Figures 4 and 5:

$$\mathbf{a} = \mu - \sqrt{\frac{\pi}{2}}\,\sigma$$

Figure 4

$$\mathbf{b} = \mu + \sqrt{\frac{\pi}{2}}\,\sigma$$

Figure 5

By taking second derivatives it can be determined that these optima are global for each of the players. These results indicate optimal bids of approximately $52.50 for the landlord and $27.50 for the tenant, which as Brams points out, is hardly a convergent outcome.

Divergence is caused by the following simple fact; each player has to make an offer that balances the increased payoff amount against the decreased probability of succeeding in the FOA. In a normal distribution there is a great deal of probability clustered near the mean, and this increased probability slopes away rapidly on each side of the mean.

It is interesting to note that these are the optimal solutions based upon the assumption that the other party is bidding in a purely rational and informed way with no risk aversion or risk seeking. If the opposite party is risk averse or unaware of the optimum bid, then the best strategy changes.

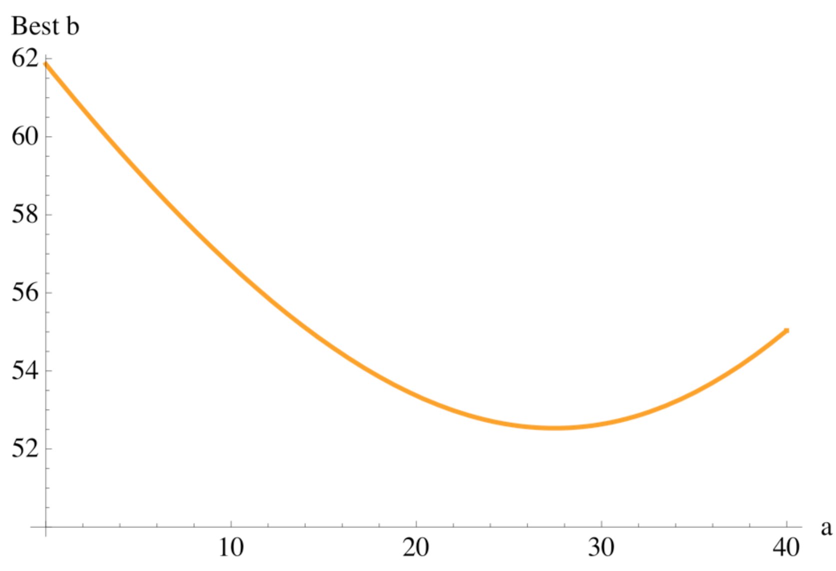

Figure 6: If a bidder knows the position of his or the opponent, there are a range of strategies. If bidder **a** chooses as in the figure (bids between 0 and the mean), **b**'s best strategy is defined on the vertical axis.

This graph clearly illustrates the tactical fact observed in the literature that a known risk-averse minimizing party should cause the maximizing party to increase his or her bid.

Description of the experiment

To explore the implications of this research on agent bidding the following behavioral economics experiment was conducted. Three groups of approximately twenty real estate brokers (all from a firm that typically represents tenants) were each asked to rapidly evaluate comparable transaction data and either submit a bid or act as the arbitrator.

The participant did three evaluations: one each in the role of arbitrator, tenant, and landlord. As well as changing roles in each round, the players also changed markets, so that their evaluation of the comparable data was *de novo* in each round.

When the participants were requested to serve in the arbitrator role, they were not asked to consider the position of the tenant and the landlord, but instead to produce a fair assessment of the market. This matches the “black box” or random variable concept of arbitrator behavior that has been widely employed in the literature. This also allows us to consider both the participants’ view of the market (when they act as an arbitrator) as well as their strategy as tenant or landlord when they serve in those roles.

The participants were told only that the exercise was designed to improve understanding of the FOA process.

Results of the experiment

Table 1 summarizes the result of the experiment.

		Market			
Role (Landlord, Arbitrator Tenant)		Chicago	Washington	New York	Orange County
A	Average submission	$30.86	$47.55	$102.03	$33.75
	Standard deviation	3.99	4.41	12.85	4.11
LL	Average submission	$33.61	$49.48	$106.40	$35.27
	Standard deviation	4.21	5.01	13.77	4.66
T	Average submission	$28.92	$43.86	$89.71	$30.43
	Standard deviation	2.91	2.88	12.32	1.24
Average across all		$31.18	$47.01	$99.23	$33.15
Standard deviation across all		4.22	4.81	14.77	4.18
Total number of submissions: 201					

Table 1

The general ordering of the results is consistent with expectations, with the average "arbitrator" position consistently between average "landlord" and average "tenant" submissions. Population standard deviations were between ten and fifteen percent of the rent—which might be considered surprisingly low.

As was the case much of the literature cited above, the bidders in the experiment consistently sub-optimized. Across the four markets the average tenant position was arbitrator mean minus .77 standard deviations, and the landlord position was arbitrator mean plus .46 standard deviations, which in either case is less aggressive than the BMM solution of mean -/+ 1.25 standard deviations.

Since the participants were directed to produce a pure estimate of market rent when functioning as the arbitrator, the mean and standard deviation of the arbitrator positions can to some extent be used as benchmarks. It is noteworthy that the landlord and tenant positions taken were generally within one standard deviation from the mean of the arbitrator position taken by the group.

One interesting result was the asymmetry between the positions taken when the participants played landlord versus tenant. As pointed out by Ashenfelter et al. (1992) a Nash equilibrium model such as Farber or Brams and Merrill should result in a bids that are symmetrical around the mean. In fact, the magnitude of difference between tenant position and arbitrator position was higher than the difference between landlord position and arbitrator position. This might be because the brokers in the experiment typically represent tenants, which might lead them to be more comfortable taking slightly more aggressive positions as a tenant than as a landlord.

However, t-tests and bootstrap analysis of the data show that the null hypothesis that the populations were drawn from the same distribution could not be ruled out at any of the standard levels of significance (p values ranged from approximately .59 to .81).

Although the results do not provide much support for symmetrical behavior, they do suggest that the bidders, whether acting in the role of landlord or tenant, were consistently risk averse. This is consistent with much of the literature, but differs from the experimental results of Ashenfelter *et al* whose experiment produced some bids that were less extreme than the anticipated risk neutral and others that were more extreme.

Risk aversion, utility theory and mean-variance analysis

Risk Aversion

Why would the players consistently choose sub-optimal strategies? There are at least two fundamental approaches in economics to modeling the propensity for individuals to be risk averse. One is based upon general utility theory; the second upon mean-variance analysis as pioneered and refined by Harry Markowitz (Markowitz (1987)).

In utility theory one assumes that every individual maximizes expected utility, subject to a function that maps the outcome of an uncertain event into a utility for the individual. These utility functions are generally speaking increasing and concave, which fits the observation that increasing amounts of wealth would logically have smaller and smaller levels of utility to an individual. A commonly employed class of utility function—the negative exponential—is depicted in Figure 7. See for example Gollier (2001) p. 27 for a review of utility functions and their behavior:

$$\mathbf{U_a}[w\_,\ \gamma\_] := -e^{-\gamma\, w}$$

Figure 7

In this function, ***w*** represents wealth and ***γ*** represents the degree of risk aversion. This function can imposed onto the BMM payoff function as in Figure 8 (see both Farber (1980) and Kilgour (1994)):

$$\mathbf{g2}[a\_,\ b\_] := \mathbf{U}[a,\ \gamma]\ \mathbf{F}\left[\frac{a+b}{2}\right] + \mathbf{U}[b,\ \gamma]\ \left(1 - \mathbf{F}\left[\frac{a+b}{2}\right]\right)$$

Figure 8

In Kilgour's 1994 article "Game-Theoretic Properties of Final-Offer Arbitration" his primary concern is to assess the effect of two different utility functions (one for each player) on the results of the game play. Our goal is to categorize the level of risk aversion necessary to cause the maximizing bidder to sub-optimize.

In order to show graphically how risk aversion changes the bidding, we normalize the initial payoff function of Figure 1 and the utility-adjusted function in Figure 8 and graph the results for parameterized version of the utility function in Figure 7. The normalization procedure is from Korsan (1994). As seen in Figure 9, a positive value of the risk aversion parameter causes the outcome of the bid to shift to the left. A consequence of this shift is that the optimal bid will be lower for the maximizing, risk averse player.

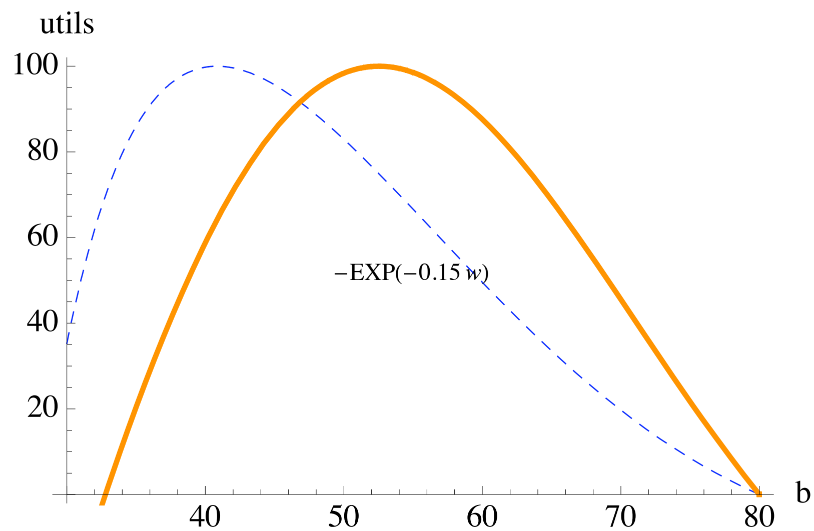

Figure 9: Employing a negative exponential utility function and a positive risk aversion parameter shifts the optimal bid for the maximizer to the left, so that a lower bid maximizes expected utility.

How much risk aversion (γ) does it take to account for the deviation from optimality seen in the experiment? For simplicity, we will examine the maximizing case only. The results are presented in Table 2. The BMM calculations are based upon the mean and standard deviation of the arbitrators in the experiment, and the results are evaluated at the optimal point (Figure 4) for the minimizing bidder.

Market	BMM "Optimal" Bid (arbitrator mean / standard deviation)	Average Experimental Bid (Landlord)	Risk Aversion Parameter Required
CHI	**36.0** (30.9/4.1)	**33.6**	**.12**
DC	**53.2** (47.6/4.5)	**49.5**	**.18**
NY	**118.7** (102.0/13/3)	**106.4**	**.07**
OC	**39.2** (33.8/4.3)	**35.3**	**.22**
		Average	**.15**

Table 2

While there is insufficient data to draw strong conclusions about the consistency of the parameters, their difference is relatively small.

How would this utility function change the overall shape of the result for the maximizing bidder across all **a** bids? To show this, we use the function in Figure 8 and the utility function in Figure 7, parameterized at approximately the average value in Table 2. The resulting graph is only meaningful for the maximizing bidder, since the minimizing bidder's utility curve will look different from the utility curve for the maximizer (for example it will likely be decreasing).

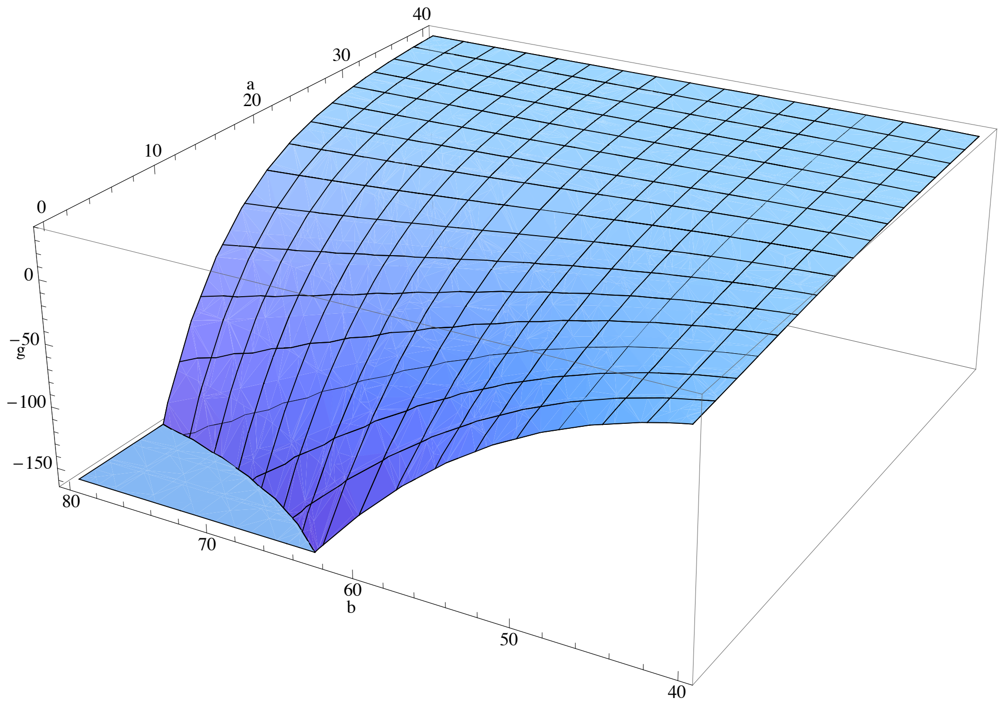

Figure 10: Shows expected outcome (g) as a function of **a** and **b** bids, assuming that the utility of the outcome is the function in Figure 7.

Figure 10 illustrates a number of interesting points. First, when the minimizing bidder chooses a bid between the mean and the optimal bid (i.e. between approximately 27.5 and 40), the bid of the maximizer is largely irrelevant when it comes to expected outcome. Second, when the maximizer chooses a bid between the mean and the optimal bid (i.e. between 40 and approximately 52.5), the outcomes are essentially similar. However, when the minimizer and the maximizer both make extreme bids, the results for the maximizer become dramatically worse. The conclusion of this graph, simply stated, is that in an environment in which the maximizer is risk averse (specifically in the case of exponential utility and with a relatively significant level of risk aversion), then the expected utility function doesn't serve as an unambiguous indicator of the appropriate game theoretical position. What it does show, however, is the dangerousness of choosing a bid that significantly exceeds the BMM maximum.

Mean-Variance Analysis

Some additional insight can be gained by looking at an issue that is not extensively addressed in the literature on FOA. What is the variance of the outcome, and how does that variance change with different plays of the game?

In the case of a discrete random variable **X**, the definition of the variance is as a summation of a product. The first term in the product is the squared difference between the discrete value (in this case **X** can either be **a** or **b**) and expected value. The second term is the probability that **X** will take on that value. Brams and Merrill provide the expected value function (g[a,b]) and the probability that either **a** or **b** will be selected, and so the variance of the BMM can be written as in Figure 11:

$$\mathtt{v}[a\_,\ b\_]\ :=\ (a - \mathtt{g}[a,\ b])^2\ \mathtt{F}\left[\frac{a+b}{2}\right] + (b - \mathtt{g}[a,\ b])^2\left(1 - \mathtt{F}\left[\frac{a+b}{2}\right]\right)$$

Figure 11

Following the ideas of mean-variance analysis (normally applied to portfolio problems) this equation makes it possible to project bids into expectation-standard deviation space. In order to graph this in two dimensions, we set the **a** bid to the theoretical optimum and graph the expected outcome of the **b** bid against the standard deviation (the square root of v[a,b] above) in Figure 12:

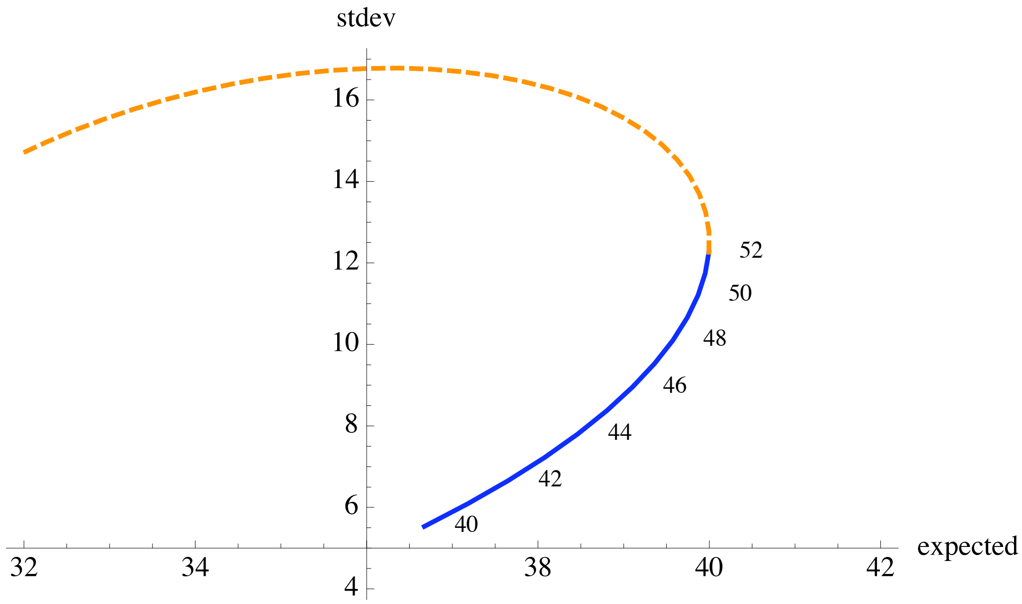

Figure 12: The expected result and its standard deviation at various levels of **b** bid. Point labels show **b** bids between the mean and the BMM optimal. All of these results produce better expected results at the same levels of standard deviation than bids above the optimal, showing that these bids dominate those above the optimal.

In Figure 12, all of the bids between the mean and the BMM optimal are rational combinations of return and risk. Choosing one requires the bidder to determine what amount of additional risk is tolerable in order to obtain a better expected outcome, and these tradeoffs can be explicitly quanitified in an agent/principal context. This range represents an efficient set, in the vocabulary of Markowitz mean-variance analysis. Bids above this point, however, are *dominated* in Markowitz terminology.

This is an important result, because it shows that a bidder sensitive to both expected outcome and standard deviation of the outcome should only consider bids between the anticipated mean and the optimal.

In the numerical example above (considering only the efficient set of bids) the standard deviation goes up 226% while the expected value increases by only 9%. This suggests that if expected value and standard deviation are considered together, a bid significantly lower than that suggested by the BMM would be selected.

Conclusion and observations

The work of Brams and Merrill and their peers make it clear that optimal play in a FOA game does not cause the bidders to converge to the arbitrator's assumed mean. However, their model does not explicitly take into account risk aversion. As several researchers have observed, risk aversion will cause maximizing parties to bid below the BMM optimal.

The empirical data in this paper supports the idea that bidders in real world settings will tend to sub-optimize. We can provide two possible explanations for sub-optimal bidding found in the empirical study supported both by the mathematical exercise and by intuition. First there is the general principle of risk aversion, which causes incremental benefits to be worth progressively less, and therefore causes the player to be unwilling to take a fair bet.

Using a widely studied utility function (the exponential), this paper computes the level of risk aversion necessary to account for the sub-optimal bidding found empirically.

The paper also presents a formula for the variance of the BMM, and makes the observation (drawn from portfolio theory) that expected outcome and the variation in the outcome relate in the following way. While a greater outcome is preferred to a lesser outcome, if two outcomes are the same, the one with lower variation is preferred to the one with greater variation. The choice among these combinations of outcome and variation is made based upon the utility function of the player. However, certain combinations of outcome and variation dominate others. This analysis shows that (for maximizers) bids above the BMM optimal are dominated by those between the mean and the BMM optimal.

One of the problems with this study (and perhaps much of economics) is that it is somewhat unsatisfying to employ mathematical models to study human behavior, when it is clear that human behavior is not algorithmic. Milton Friedman addressed this problem cogently in his seminal essay “The Methodology of Positive Economics” (in several books, including Friedman (1953)).

In that essay Friedman makes the point that positive economics should be judged in terms of the results of its predictions, not based upon whether the underlying model or analysis precisely maps the actions of those involved. What positive insights can this analysis provide for practitioners?

1. When evaluating the actions of an arbitrator, one should consider both the likely outcome and the variability of that outcome. The brokers in this study computed means of comparable transactions and normalized current and future payments pretty effectively. To go one step further and compute the variability intrinsic in market data is a logical step.

2. Armed with public information on comparables, and with a computation of an arbitrator’s expected value and the likely standard deviation around that mean, it becomes possible to quantify the FOA using the BMM model. This can be done easily in Excel and can be used to calibrate bids and expectations of results.

3. The BMM model states that risk neutral bidding will in general produce better results. However, these improved results come with a price in risk. The mean-variance analysis in

this paper allows the tradeoffs among different strategies to be depicted and better understood. The plot presented in Figure 12 can be used both by agents to understand the tradeoffs in their recommend bids, and principals in understanding the consequences of FOA versus negotiated settlement.

4. If it is possible to understand (in advance) what the counter party's attitude toward risk might be, a risk-neutral practitioner can significantly improve their position. Figure 6 shows that foreknowledge of the other party's risk aversion can be intelligently used to increase the aggressiveness of one's bid.

5. It is possible that the benefits of this approach will over time be minimized in classic prisoner's dilemma fashion as found in Ashenfelter and Dahl (2005). Until that happens, the more numerate practitioner has an advantage.

Returning to the original questions discussed in the introduction, we conclude that there is a certain amount of truth to each of the described perspectives on FOA in real estate. While Farber, Brams and Merrill, and their fellow researchers show that FOA is not mathematically convergent, considerations of variability make it clear that it is not unreasonable to bid somewhere between the anticipated mean and the BMM optimal position.

Therefore, adopting positioning between theoretical optimum and "safe" but sub-optimal bidding is supported both by analyzing the problem through the lens of traditional utility theory and the lens of mean-variance analysis.

It could be argued that this conclusion supports the idea that FOA drives the parties to be more reasonable, even if the process is not technically convergent. There should be no question that the outcome frontier depicted in Figure 10 should give any bidder pause when considering an extreme bid. That, when combined with Hanany et al. (2007)—who demonstrate an optimal outcome can be achieved when bargaining precedes FOA—helps to explain why so many FOA negotiations, particularly in real estate, get settled prior to the application of FOA.

References

Ashenfelter, Orley, and David E. Bloom. 1984. Models of Arbitrator Behavior: Theory and Evidence *The American Economic Review,* **74** (1) 111-124.

Ashenfelter, Orley, and Gordon B. Dahl. 2005. Strategic Bargaining Behavior, Self-Serving Biases, and the Role of Expert Agents: An Empirical Study of Final-Offer Arbitration *NBER Working Paper 11189.*

Ashenfelter, Orley, Gordon B. Dahl, Janet Currie, Henry S. Farber, and Matthew Spiegel. 1992. An Experimental Comparison of Dispute Rates in Alternative Arbitration Systems *Econometrica,* **60** (6) 1407-1433.

Brams, Steven J. 2003. *Negotiation Games.* Routledge, New York.

Brams, Steven J., and Samuel Merrill. 1983. Equilibrium strategies for final-offer arbitration: there is no median convergence *Management Science,* **29** (August) 927-941.

Brams, Steven J., and Samuel Merrill. 1991. Final Offer Arbitration with a Bonus *European Journal of Political Economy,* **7** 79-92.

Dickinson, David L.. 2004. A Comparison of Conventional, Final-Offer, and "Combined" Arbitration for Dispute Resolution *Industrial & Labor Relations Review,* **57** (2) 288-301.

Farber, Henry S.. 1980. An Analysis of Final-Offer Arbitration *Journal of Conflict Resolution,* **24** (4) 683-705.

Faurot, David J., and Stephen McAllister. 1992. Salary Arbitration and Pre-Arbitration Negotiation in Major League Baseball *Industrial & Labor Relations Review,* **45** 697-710.

Friedman, M. 1953. In *Essays in Positive Economics*. Chicago: University of Chicago Press.

Gollier, Christian 2001. *The Economics of Risk and Time.* The MIT Press, Cambridge, Massachusetts.

Hanany, Eran, D. Marc Kilgour, and Yigal Gerchak. 2007. Final-Offer Arbitration and Risk Aversion in Bargaining *Management Science,* **53** (November) 1785-1792.

Kilgour, D. Marc. 1994. Game-theoretic properties of Final-Offer Arbitration *Group Decision and Negotiation,* **3** (3) 285-301.

Korsan, Robert J.. 1994. Nothing Ventured, Nothing Gained: Modeling Venture Capital Decisions *The Mathematica Journal,* **4** (1) 74-80.

Marburger, Daniel R., and John F. Scoggins. 1996. Risk and Final Offer Arbitration Usage Rates: Evidence from Major League Baseball *Journal of Labor Research,* **XVII** (4) 735-745.

Markowitz, Harry M. 1987. *Mean-Variance Analysis in Portfolio Choice and Capital Markets.* Basil Blackwell, New York.