The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXVIII, NO. 5 • OCTOBER 2013

Trading Complex Assets

BRUCE IAN CARLIN, SHIMON KOGAN, and RICHARD LOWERY*

ABSTRACT

We perform an experimental study to assess the effect of complexity on asset trading. We find that higher complexity leads to increased price volatility, lower liquidity, and decreased trade efficiency especially when repeated bargaining takes place. However, the channel through which complexity acts is not simply due to the added noise induced by estimation error. Rather, complexity alters the bidding strategies used by traders, making them less inclined to trade, even when we control for estimation error across treatments. As such, it appears that adverse selection plays an important role in explaining the trading abnormalities caused by complexity.

COMPLEXITY BOUNDS THE ABILITY of market participants to accurately value assets. Some assets are easier to analyze (e.g., Treasury bonds), whereas others have unbounded contingencies that prevent agents from pinning down their exact values (e.g., corporate bonds with embedded American options and credit default swaps). Indeed, for many financial assets there do not exist closed-form analytical solutions to quantify their value. Moreover, the recent growth of securitization and the use of finer tranches of underlying collateral are likely to exacerbate complexity in the market (e.g., Furfine (2011)).

The purpose of this paper is to study how complexity affects asset trading, price volatility, liquidity, and trade efficiency. Surely complexity increases uncertainty regarding asset values, but it remains unclear what effect this should have on trading and why. Especially since complexity makes it harder for traders to know whether counterparties have an informational or skill advantage, there are multiple potential channels through which complexity may affect asset pricing and liquidity.

We study complexity in a laboratory setting. Participants were asked to evaluate the price of certain assets and were then given the opportunity to trade based on their information. Each subject participated in 15 distinct periods, each of which consisted of two stages. In the first stage, each participant was given information regarding several portfolios composed of four assets and was

*Carlin is from the Anderson School of Management, University of California, Los Angeles. Kogan and Lowery are from the McCombs School of Business, University of Texas at Austin. We would like to thank Kim Donghyun for excellent research assistance and finance seminar participants at the Australian National University, Hebrew University, Melbourne University, Michigan State University, University of New South Wales, University of Queensland, University of Sydney, Tel Aviv University, University of Texas at Austin, and the 2010 Miami Behavioral Finance Conference for their comments and suggestions.

DOI: 10.1111/jofi.12029

asked to submit his best estimate of the value of a particular asset included in these portfolios. Following that, in the second stage, participants were randomly paired and were given the opportunity to trade the asset through a well-defined anonymous bargaining process. Each pair contained one buyer and one seller such that the buyer's private value of the asset was always greater than the seller's private value of the asset. The complexity involved in assessing the asset's value from the portfolios varied across periods, and we collected information regarding frequency of trade, trading prices, and trading surplus as a function of complexity.

Not surprisingly, complexity increased estimation error and price volatility. However, higher price volatility did not appear to be simply due to the added noise induced by increased estimation error. First, we estimated the sellers' and buyers' bidding strategies as a function of their *guess estimates* with and without complexity. After controlling for estimation error, we found that the bid as a function of each trader's estimate was much flatter when the portfolio problem was more complex. Second, we decomposed the price volatility into guess deviations (distance between the guess and the asset's value) and bid deviations (distance between the bid and the guess) and found that bid deviations were almost twice as large in the complex treatments.

Complexity also affected the liquidity of the assets under study and the resultant trade efficiency, as measured by the trade surplus generated in each period. Making the required computation simpler increased trade efficiency by 11% (from 73% to 81%); this was more pronounced when more bidding rounds were allowed. Whereas efficiency rose from 73% to 84% for the simple treatment when the number of rounds increased from one to three, efficiency remained unimproved in the complex treatment (72% vs. 73%). Overall, aggregate surplus tended to be lower in the complex cases.

To investigate our findings further, we ran a second set of experiments in which participants faced uncertainty about their private values but did not have to solve portfolio problems. We linked these experiments to the previous ones by giving subjects noisy signals about their private value, where we added uncertainty according to the observed average estimation error from the first experiments. We found no statistically discernible difference in efficiency, liquidity, or price volatility in these experiments when we varied the uncertainty subjects faced about their private values. This provides convincing evidence that added noise due to estimation error is unlikely to account for the trading effects induced by complexity.

Why, then, does complexity induce these effects? The answer probably lies with adverse selection. As problem solving becomes more challenging, people become concerned about whether their trading partner knows more. Indeed, disparity in financial expertise has been shown to be destabilizing because it worsens adverse selection problems in trading venues (e.g., Glode, Green, and Lowery (2012)). In our experiments, this is supported by the fact that bidding becomes more conservative in the complex treatments. However, as we also document, subjects were more reluctant to trade at all in the complex treatments. We identified cases in which extreme bids were offered that had no chance of inducing a loss (e.g., a seller offering an ask greater than or equal to the highest possible asset value). Such offers were almost four times more likely in complex periods than in simple ones.

Our experimental findings can be justified theoretically. Without any uncertainty over private values, the bargaining game used in our experiment gives rise to an equilibrium that does not achieve full efficiency. Adding exogenous noise may not exacerbate the problem; consider, for example, an extreme case in which the problem is so complex that it provides the players with no information about their private values. In such a case, full efficiency can be supported in equilibrium as neither player's bid is a function of his private value and each player knows that his counterparty is also uninformed. However, in the case in which one trader has not figured out the problem and the other has (or at least there is a belief over whether they did), an adverse selection problem may limit trade. Exploring the underpinnings of this interaction is the focus of future research.

The subjects in our study viewed valuation requiring computational difficulty as more complex: they were more reluctant to trade, even if they correctly determined the value of the asset. As such, there was lower trade efficiency in complex treatments. However, it is important to note that computational difficulty is only one potential proxy for asset complexity and may not be a sufficient condition for particular securities to be *perceived* as complex. For example, consider assessment of the value of Goldman Sachs stock. Its portfolio is immense, changes dynamically, and has unbounded contingencies. Yet many investors might not perceive Goldman Sachs stock to be complex, especially when compared to an exotic option or a serially securitized asset.[1] In this case, investors might substitute diversification as a proxy for lack of complexity. That is, since there are diversification benefits offered by holding Goldman Sachs stock, traders may not perceive it to be complex, though it clearly is from a computational standpoint. Therefore, while we use computational complexity as a convenient experimental proxy in the analysis here, we do not assert that this is the only way in which complexity may arise or be perceived.

As pointed out by Brunnermeier and Oehmke (2010), asset complexity may have asset pricing implications and may drive how assets are managed and traded. For example, Arora et al. (2010) show that, once complexity is taken into consideration, derivatives can actually increase asymmetric information costs instead of decreasing them. Bernardo and Cornell (1997) provide empirical evidence that the complexity of collateralized mortgage obligations causes the variance of bids to be much larger than can be explained by estimation error alone. Furfine (2011) studies commercial mortgage-backed security deals and shows that postcrisis loan performance was worse for loans packaged in more complex securitizations, even though investors bought such securities at a premium compared to less complex deals. In comparison to previous empirical studies, our experimental investigation does not use professional traders from real markets. However, our study complements previous work because we

[1] The authors thank Pierre Collin-Dufresne for pointing this out and for this example.

control for confounding variables that would make real-world tests challenging: imperfect information, hidden attributes (e.g., quality), relationships between traders, self-selection, and the innate liquidity of assets. Moreover, our work makes several novel predictions that might be studied in financial markets. For example, our results imply that regulation requiring asset standardization should decrease price volatility, increase liquidity, and generate welfare (though increased trade surplus). Likewise, assets with more contingencies should have more price volatility than predicted by the underlying assets used to replicate them.

Our work adds to a growing literature on complexity in financial markets, which demonstrates that complexity is a robust concern that is not alleviated by competition. Carlin (2009) studies the effect of competition on complexity and shows that, as the number of firms rises, each firm adds more complexity to its prices. Carlin and Manso (2011) show that educational initiatives undertaken by a social planner to increase sophistication may worsen the amount of complexity in the market. Payzan-LeNestour and Bossaerts (2011) study complexity in a laboratory setting and show that subjects often use Bayesian learning when they face such problems.

However, to our knowledge, our paper is the first to explore the effect of complexity on asset trading. Healy et al. (2010) examine the performance of different market mechanisms in aggregating information in simple or complex environments, but their setting does not permit direct comparisons of the performance of traders in these different settings. Other experimental papers also explore various aspects of bargaining games related to ours, which is often referred to as a sealed-bid double auction with incomplete information. Chatterjee and Samuelson (1983) theoretically analyze such a bargaining game and show that the Nash equilibrium strategy is monotonic in bidders' reservation values. Radner and Schotter (1989) test this experimentally and find that subjects do use strategies that approximate monotonic linear bidding functions, and that subjects capture a large fraction of the available trading surplus. Schotter (1990) discusses a large set of experiments using the same bargaining mechanism while varying different features of the environment. Bidding strategies remain largely monotonic, if not always linear, but the efficiency of the mechanism remains intact. Our work adds to theirs in several respects. Whereas Radner and Schotter (1989) and Schotter (1990) provide their subjects with precise information about their private values, we do not. Instead, our subjects are given full information in a form that requires computation. In some cases, this may lead subjects to have uncertainty about their private value, which may affect their trading behavior. Indeed, as we show, complexity can affect the linear bidding strategy, making it less responsive to changes in value estimates. Further, as we show, while the simple treatments with three bidding rounds reproduced the trade efficiency in Schotter (1990), complexity reduced it.

Finally, Radner and Schotter (1989) and Schotter (1990) focus on cases in which traders' private values are independent. In contrast, we consider the case in which private values are affiliated, which is more relevant to the analysis of financial markets. In this sense the closest theoretical work to ours is Kadan (2007), who investigates the theoretical properties of k-double auctions (a generalization of the bargaining mechanism we employ) with affiliated private values. Kadan (2007), however, does not consider the case in which traders are not directly given their value and also does not consider the specific private value model we consider.

The rest of the paper is organized as follows. In Section I, we describe our experimental setup. In Section II, we describe our data. Section III characterizes our results. Section IV provides concluding remarks. The Internet Appendix available in the online version of this article includes the instructions that subjects received during the experiments and results from robustness checks that we discuss in Section III.E.

I. Experimental Design

In this study, we ran two experiments. In both experiments, subjects with private values were given opportunities to trade. Experiment 1 required subjects to ascertain their private value by solving portfolio problems. Experiment 2 did not require such computation. Instead, subjects received noisy signals about their private values. Every subject in the study participated in one and only one experimental session. No subject participated in both Experiment 1 and Experiment 2.

A. Experiment 1: Computational Complexity

Each experimental session consisted of 15 periods. At the beginning of each period, every subject was given information about the composition and value of four portfolios. After studying this information, subjects were asked to estimate the value of a traded security contained in the portfolios, and this estimate was recorded. Each subject was then allowed to trade with an anonymous partner (i.e., another subject) in a well-defined simple bargaining process that we specify shortly. Assets were traded in Experimental Currency Units (ECUs), with one ECU equal to 10 cents.

The value of the particular security of interest could be solved deductively by using the principal of no arbitrage. Specifically, subjects received information about four baskets of securities, labeled Basket 1 through Basket 4. Each basket contained quantities of four securities, labeled Security A through Security D. Subjects were given information such as the number of units of each security in each basket and the price of each basket. Figures 1 and 2 provide examples of typical problems that a subject might face.

Given the information provided, the problems faced by subjects were either Simple (as in Figure 1) or Complex (as in Figure 2). Simple portfolio problems could often be solved by inspection, or with minimal computation. Complex problems required more effort and ingenuity. However, no matter how challenging the problem, the information was sufficient to determine the price of

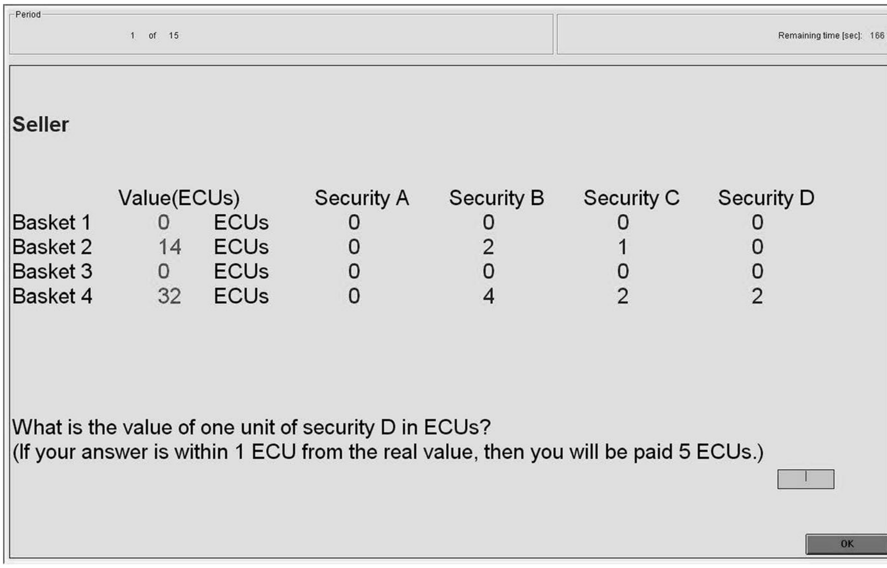

Figure 1. Screenshot example—simple condition. This is a screenshot from the interface used for the experiment. It provides an example of a decision problem used in the simple condition.

the traded security with certainty. Each subject was given three minutes to estimate the asset's value (i.e., assess the value of Security D).

We divided the 15 periods in the session into three sets of five periods, with each set containing either simple or complex problems. We made sure that every 15-period session included at least one set of simple and one set of complex periods, so that we could study within-subject variation in behavior. In each period, subjects were asked to submit their best estimate of the fundamental value of the security in question. Subjects whose guess fell within one unit of their true private value received an additional five ECUs.[2]

In each period, the actual private value of the asset under consideration depended on whether the subject was assigned to be a buyer or a seller. The seller's private value was drawn from a uniform distribution from 1 to 20. The buyer's private value was equal to the seller's realized private value plus an additional margin that was drawn from a uniform distribution from zero to 20. The two random draws were independent across subjects and periods. Therefore, the seller's value ranged from 1 to 20, whereas the buyer's value ranged from 1 to 40. Due to this construction, the buyer's value was somewhat more informative. For instance, if the buyer valued the asset at 40, she knew that the seller had a private value of 20 for sure. Alternatively, if the buyer assessed a private value of 20, she would not be able to update her beliefs about

[2] The first two sessions were conducted without explicitly rewarding subjects for accurate guesses. However, we find no difference in subjects' errors between the first two sessions and the remaining sessions.

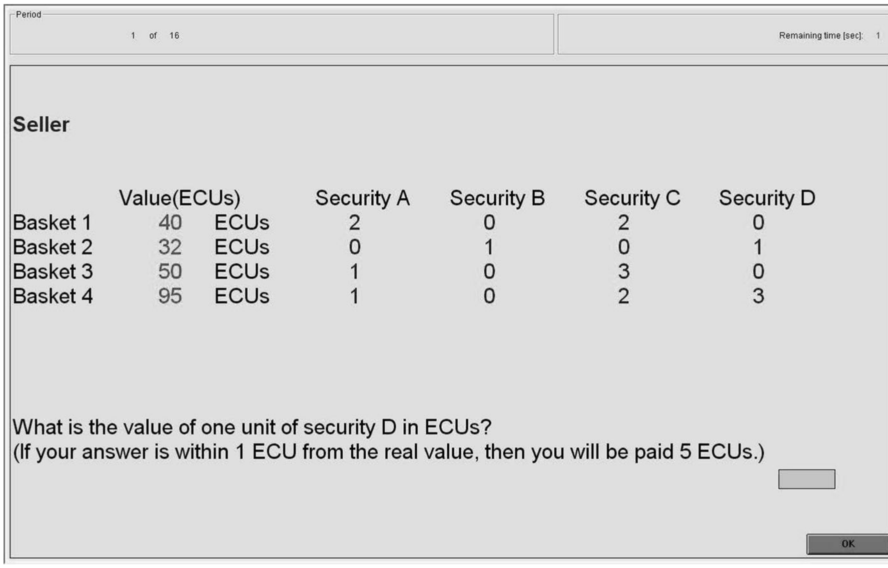

Figure 2. Screenshot example—complex condition. This is a screenshot from the interface used for the experiment. It provides an example of a decision problem used in the complex condition.

the seller's private value beyond her prior that it was uniformly distributed between 1 and 20.

We chose this structure to ensure that there would be positive gains in every period, that private values were affiliated,[3] that the supports for the parties were overlapping, and that the design would be easy to explain to subjects. Assuring positive gains allows us to compare the likelihood of trade in the different treatments to a baseline in which trade is always efficient. Making the values affiliated makes our analysis relevant to finance. That is, whereas bargaining with independent private values is more common in the experimental literature, we view affiliation as more appropriate when considering financial instruments. Having overlapping supports is necessary to ensure that traders would not ignore their signals when trading. Indeed, if supports were not overlapping, traders might simply trade even if they could not determine their private value. Finally, subjects easily understood how the private values were determined.

After submitting their best estimate of their private values, subjects participated in a bargaining game. They were assigned the role of a buyer or a seller, and were randomly paired with another subject with the opposite role. Subjects were allowed to trade anonymously over either one or three bargaining rounds chosen randomly and announced at the beginning of each period. For any bargaining round, the subjects were given 30 seconds to simultaneously submit

[3] Recall that affiliation means that a higher value for the buyer will generally be associated with a higher value for the seller, and vice versa.

bids. If either trading partner failed to do so, the bidding round is terminated. If the bid submitted by the buyer (weakly) exceeded that submitted by the seller, a transaction occurred in which the buyer paid the seller the average value of the bids. The payoff from a trade for the buyer was equal to their private value minus the traded price. Likewise, the payoff from a trade for the seller was equal to the transaction price minus his private value. In periods with one bargaining round, if no transaction took place, no more bargaining was allowed. In periods with three bargaining rounds, if a bargaining round did not result in a transaction, subjects were notified that no transaction occurred and they progressed to the next round without being informed of each others' bids. If a transaction did not occur by the end of the third bargaining round, subjects forfeited any value from trade.

The following summarizes what information the subjects had during each session. First, before each subject began the study, they were given full instructions regarding the protocol. We confirmed understanding of the instructions by giving each subject a formal quiz to test his proficiency regarding the protocol.[4] Following that, at the time that subjects were given each portfolio problem, they were also assigned a role of seller or buyer, and were told whether there would be one or three bargaining rounds in the trading game. During each session, we did not allow subjects to communicate with each other. Information collected from subjects and trade between them occurred anonymously via a computer terminal, using a standard z-tree program (Fischbacher (2007)). Finally, it is important to note that both buyers and sellers were provided the same set of assets in the portfolios, with the only difference being the price of the portfolios. Thus, the buyer and the seller in each interaction faced the same level of difficulty in ascertaining their own private value and were made aware of that.

At the end of each period, subjects were informed whether a trade had occurred and the value of Security D to them. At the end of the experiment, subjects received a detailed account of their ECU earnings in each period, and their total pay, which was remitted in U.S. dollars.

B. Experiment 2: Private Value Uncertainty

After observing our results for Experiment 1, we ran a second experiment to directly test whether the treatment results obtained from complexity are solely due to changes in private value uncertainty. In Experiment 2, the subjects did not solve portfolio problems but faced uncertainty regarding their private values. This experiment followed the same protocol as the first with one exception: instead of providing subjects with information that they could use to deduce their private value, we provided them with a noisy signal about his private value.

Each experimental session consisted of 15 periods. At the beginning of each period, each subject received a clue about his private value. They were

[4] Screen shots with instructions for Experiment 1 are available in the Internet Appendix.

informed that with some probability the clue is correct (i.e., the private value is equal to the clue) and with the complementary probability the clue is simply noise (i.e., a signal drawn from the same distributions used in the first experiment).

The uncertainty we imposed in Experiment 2 was designed to mimic the estimation errors observed in Experiment 1. Specifically, we measured the private value uncertainty for each of the 20 possible problems presented to subjects in Experiment 1. For each problem, we calculated the fraction of times that subjects guessed their private value correctly (i.e., their guess was at most one unit away from the true value). We then used these probabilities when drawing clues for subjects in the second experiment. In short, we designed the second experiment so that the level of uncertainty over asset values was the same as in the first experiment.

We divided the 15 periods in each session into three sets of five periods, with each set containing either low-uncertainty or high-uncertainty treatments. As in Experiment 1, we made sure that every 15-period session included at least one set of low and one set of high uncertainty periods, so that we could study within-subject variation in trading behavior.

During each period, after each subject received his clue, he was assigned the role of a buyer or a seller, and were randomly paired with another subject with the opposite role. As in Experiment 1, subjects were allowed to trade anonymously over either one or three bargaining rounds, chosen randomly and announced at the beginning of each period. During each session, we did not allow subjects to communicate with each other. As in Experiment 1, the subjects were given 30 seconds to simultaneously submit bids. If either trading partner failed to do so, the bidding round terminated. If the bid submitted by the buyer (weakly) exceeded that submitted by the seller, a transaction occurred in which the buyer paid the seller the average value of the bids. The payoffs to the buyer and seller were calculated as before.

As in Experiment 1, before each subject began the study, he was given full instructions regarding the protocol and we confirmed understanding of the instructions with a quiz.[5] It is important to note that both buyers and sellers were provided with the same uncertainty in each period, with the only difference being the price of the asset in question. Finally, at the end of each period, subjects were informed whether a trade had occurred and the value of the security to them. At the end of the experiment, subjects received a detailed account of their ECU earnings in each period and their total pay, which was remitted in U.S. dollars.

II. Data

The data were collected over the course of 10 independent sessions at the McCombs School of Business at the University of Texas at Austin. Sessions typically lasted just over an hour. The average pay in Experiment 1 was $15

[5] Screen shots with instructions for Experiment 2 are available in the Internet Appendix.

(including the $5 show-up fee), with a standard deviation of $4.30. The average pay in Experiment 2 was $11, with a standard deviation of $5.20.

Table I describes the data collected. In total, 70 subjects participated in Experiment 1 and 48 subjects participated in Experiment 2. No subject attended more than one session and no subject participated in both experiments. In Experiment 1, there were 242 periods in which subjects received the simple treatment and 272 in which they received the complex treatment. In Experiment 2, there were 219 periods in which subjects received the low-uncertainty treatment and 140 in which they received the high-uncertainty treatment.

At the end of each session, subjects completed a short demographics survey. The characteristics of the subjects in each experiment were quite similar.[6] In our sample of subjects in Experiment 1, the majority were male (61 of 70), majored in economics (47 of 70), and were at an advanced stage of their school work (47 were third year or higher). The mean (self-reported) grade point average (GPA) was 3.48. In our sample of subjects in Experiment 2, the majority were again male (44 of 48), majored in economics (37 of 48), and were at an advanced stage of their school work (31 were third year or higher). The mean (self-reported) GPA was 3.44.

We did not apply any filtering to the data. The only observations that were dropped from the analysis were rounds in which one or both subjects did not submit a bid. During Experiment 1, this happened in 33 rounds (14 of which occurred in the simple condition and 19 in the complex condition) of the 804 total rounds of the experiment. During Experiment 2, this only happened in 9 rounds out of a total of 505 rounds.

III. Analysis

A. Estimation Errors

In Experiment 1, we investigate the effect of complexity on estimation errors by analyzing the data at the session level. We report all of the variables of interest as an average across both subjects and periods and calculate the *p*-values from a *t*-test of mean differences between the two conditions (simple vs. complex). This makes the most conservative use of the data as it treats all observations collected in a given session and treatment condition as a *single observation*. This design is intended to capture any correlations across periods and subjects.

Panel A of Table II confirms that periods designed to be complex were perceived by subjects to be different from periods designed to be simple. Comparing the estimates reported by subjects with the actual value of the security to assess how estimation errors varied with treatments, it is clear that complexity leads to more estimation errors. Subjects guessed the value correctly 72% of

[6] As we will show in Section III, the average levels of the dependent variables that we study are generally comparable across the two experiments, which suggests that our results are unlikely to arise just due to variation between the subjects who participated in the first and second experiments.

Table I

Experiments 1 and 2: Data Summary

The table reports the data collected in the experiment, divided into sessions. It reports (in the order of the columns) the number of subjects per session, the periods in which the simple condition was conducted, the periods in which the complex condition was conducted, the total number of period observations, the total number of round observations, the average payoff (across subjects and periods), the number of period observations collected under the simple condition, and the number of period observations collected under the complex condition.

Panel A: Complexity Sessions

Session ID	*N* of subjects	Simple periods	Complex periods	*N* of periods	*N* of rounds	Average payoff	*N* of simple periods	*N* complex periods
1	16	6–10	1–5, 11–15	115	166	3.93	40	75
2	8	11–15	1–10	60	87	3.87	20	40
3	8	1–5, 11–15	6–10	60	105	3.69	40	20
4	18	6–10	1–5, 11–15	133	190	3.83	45	88
5	20	1–5, 11–15	6–10	146	223	3.57	97	49
Total	70	NA	NA	514	771	3.78	242	272

Panel B: Uncertainty Sessions

Session ID	*N* of subjects	Low uncertainty periods	High uncertainty periods	*N* of periods	*N* of rounds	Average payoff	*N* of low uncertainty periods	*N* of high uncertainty periods
11	12	1–5, 11–15	6–10	90	143	4.5	60	30
12	8	6–10	1–5, 11–15	59	67	3.36	19	40
13	10	1–10	11–15	75	105	3.67	50	25
14	10	1–5, 11–15	6–10	75	103	3.23	50	25
15	8	1–5, 11–15	6–10	60	78	3.10	40	20
Total	48	NA	NA	359	496	3.57	219	140

Table II

Experiments 1 and 2: Result Summary—Complexity and Uncertainty

The table reports the average level of various measures across the two treatment conditions (simple and complex) and sessions. Panel A focuses on measures of complexity: the fraction of periods in which subjects provided an exact guess of their private value, the average guess error, and the fraction of rounds in which buyers (sellers) submitted bids that were higher (lower) then their maximum (minimum) possible value. Panel B focuses on the main dependent variables: the average fraction of surplus captured by both subjects, the fraction of periods resulting in a transaction, the average payoff asymmetry (across the two subjects), the average bid gap (in rounds that did not result in a transaction), the average number of rounds used before a transaction took place, and the adjusted price volatility. Panel C reports the main dependent variables for the uncertainty sessions. The final column reports the p-values from a mean equality test across treatments, treating each session as an observation.

		Session Number						
		1	2	3	4	5	Total	p-value
Panel A: Measures of Complexity								
Frequency of exact guesses	Simple	0.69	0.88	0.84	0.67	0.68	0.72	0.002
	Complex	0.25	0.35	0.57	0.38	0.16	0.31	
Average guess error	Simple	3.14	1.45	1.01	2.92	2.57	2.38	0.002
	Complex	9.75	7.39	5.20	7.35	5.65	7.55	
Frequency of bid error	Simple	0.018	0.017	0.000	0.018	0.003	0.008	0.146
	Complex	0.018	0.044	0.000	0.022	0.016	0.022	
Panel B: Dependent Variables								
Average efficiency	Simple	0.88	0.68	0.78	0.84	0.79	0.81	0.020
	Complex	0.80	0.68	0.71	0.73	0.67	0.73	
Frequency of transaction	Simple	0.85	0.55	0.62	0.71	0.66	0.69	0.169
	Complex	0.72	0.60	0.55	0.69	0.57	0.65	
Payoff asymmetry	Simple	5.49	2.64	4.70	5.83	3.34	4.42	0.033
	Complex	7.13	5.42	4.95	6.57	4.84	6.21	
Bid gap	Simple	2.67	3.06	3.50	5.27	3.03	3.48	0.060
	Complex	3.86	3.94	3.94	5.31	3.24	4.16	
Number of rounds used	Simple	1.62	1.14	1.67	1.44	1.39	1.47	0.433
	Complex	1.52	1.00	1.33	1.72	1.25	1.49	
Price volatility	Simple	7.26	3.89	6.86	10.31	4.70	6.99	0.041
	Complex	9.75	9.04	7.91	10.77	7.97	9.70	

		Session Number						
		11	12	13	14	15	Total	p-value
Panel C: Dependent Variables—Uncertainty Sessions								
Average efficiency	Low uncertainty	0.79	0.68	0.77	0.69	0.67	0.74	0.886
	High uncertainty	0.87	0.73	0.47	0.72	0.76	0.72	
Frequency of transaction	Low uncertainty	0.72	0.58	0.70	0.62	0.62	0.66	0.794
	High uncertainty	0.87	0.65	0.44	0.68	0.70	0.67	
Payoff asymmetry	Low uncertainty	9.57	4.64	4.73	5.21	6.52	6.57	0.552
	High uncertainty	6.35	6.13	6.09	9.24	6.68	6.83	
Bid gap	Low uncertainty	4.50	3.50	3.20	3.39	3.83	3.71	0.597
	High uncertainty	5.38	3.36	2.71	3.00	4.92	3.48	
Number of rounds used	Low uncertainty	1.52	1.00	1.75	1.47	1.07	1.43	0.069
	High uncertainty	1.54	1.33	2.25	1.58	1.90	1.65	
Price volatility	Low uncertainty	9.90	6.74	5.92	5.83	8.26	7.98	0.411
	High uncertainty	7.75	8.07	7.39	9.84	8.23	8.41	

the time in the simple treatment and 31% of the time in the complex treatment. Likewise, the average guess error—the absolute difference between the guess and the private value—increased from 2.38 in the simple treatment to 7.55 in the complex treatment. As the mean test results at the session level suggest, these differences are all statistically significant at the 5% level. Finally, we study bid errors, not just estimation errors. While it is hard to determine what a suboptimal strategy might be, it is clear that buyers cannot benefit from submitting bids that exceed their highest possible private value. Likewise, sellers cannot benefit from submitting bids that are lower than their lowest possible private value. In both treatments, the frequency of bid errors is low (2.2% vs. 0.8%), but the difference is not statistically significant at the 10% level.

B. Trading Effects

Since our goal is to determine the effect of complexity on price volatility, liquidity, and trade efficiency, we construct the following dependent variables:

- (i) Price Volatility: the volatility of normalized transaction prices across groups and periods. Normalized transaction prices were obtained by scaling prices by the midpoint of subjects' private values in the period. This normalization controls for the contribution random draws of private values might have on observed transaction prices.
- (ii) Frequency of Transaction: a per-period proportion of the total number of interactions in which a transaction took place.
- (iii) Bid Gap: the difference between the buyer's bid and the seller's bid in the last round of bidding in periods in which traders failed to transact. This measures how close the traders were from coming to agreement, and provides another measure of liquidity.
- (iv) Number of Rounds Used: the number of rounds used before a transaction took place, conditional on a period allowing three rounds of bidding. Three-round periods that did not result in a transaction were excluded since these periods have three rounds by definition.[7]
- (v) Efficiency: the proportion of the total payoff from trading obtained by the subjects across a number of periods divided by the total surplus from trade available in these periods. Surplus is measured by the difference between the buyer's and seller's private values. In contrast to frequency of transaction, efficiency is affected by the surplus available in periods in which no transaction took place.[8]

[7] Our results are qualitatively unchanged when we define the measure without excluding three-round periods that did not result in a transaction.

[8] Note that, on a transaction level, frequency of transaction and efficiency are identical. Both are equal to one if a transaction occurs and zero otherwise. However, summing aggregate surplus generated and dividing by the total surplus available across periods, implies that the two measures will yield different values.

Using these variables, we employ a within-subject experiment design. Since our focus was on the effect of complexity, we treat every session as a unit of observation, averaging the dependent variable (e.g., frequency of transaction) within each of the conditions and taking the difference between the two averages. Therefore, each observation in the tests that we performed is the difference in the level of the dependent variable across the two treatments for a given session. Under the null, this difference should be zero. This approach has some straightforward advantages. First, it controls for the idiosyncratic attributes of each subject's individual behavior as it is netted out when taking differences. Second, it controls for the correlation in behavior across periods for a given subject by measuring the dependent variable as the average across all periods.

B.1. Price Volatility

Panel B of Table II demonstrates that complexity increased price volatility by 38% in Experiment 1. This difference is economically large and statistically significant ($p = 0.04$). To identify whether this increase was due to the noise added by higher estimation error or from a change in trading behavior, we decompose the price volatility into two components: guess deviation and bid deviation.[9] Recall that we measure price volatility by computing the standard deviation of normalized transaction prices across periods, where normalized transaction prices are equal to the difference between the average of the two parties' bids. As such, transaction prices may deviate from a private value midpoint due to a deviation of bids from estimates, and/or a deviation of estimates from private values. We denote the first deviation as a "bid deviation" and the second as a "guess deviation." We compute these measures of volatility across sessions and treatment conditions. Consistent with the idea that volatility is driven by private value estimation errors, we find that guess deviations are on average almost twice as large in the complex treatment compared with the simple treatment (7.9 vs. 3.9). However, and equally important, we also find that bid deviations are significantly larger under complexity (7.3 vs. 4.3). This suggests that complexity induces changes in *bid behavior* and not just estimation errors.

Results from Experiment 2 support this conclusion. Panel C in Table II summarizes the results for the same set of dependent variables that we studied in the complexity treatment. In the panel, the first row under each dependent variable corresponds to data collected from low-uncertainty periods, while the second row under each dependent variable corresponds to data collected from high-uncertainty periods. According to Panel C, there is no significant difference in price volatility between low- and high-uncertainty treatments, which

[9] In our computations, we also took the correlation between these two variables into consideration. However, we did not find this correlation to be statistically different across treatments. These calculations are available from the authors upon request.

provides direct evidence that price volatility is not driven by the level of uncertainty over private values.

B.2. Liquidity and Trade Efficiency

Panel B of Table II also demonstrates how complexity affected our liquidity measures in Experiment 1. We find that complexity decreased efficiency by 10% and increased the bid gap by 20%. These differences are economically large and statistically different from zero at the 10% significance level. Interestingly, uncertainty does not affect these measures in Experiment 2. According to Panel C, we find no statistically or economically discernible difference in efficiency or liquidity. The only statistically significant difference is the number of rounds used—it is somewhat higher under high levels of uncertainty. Interestingly, this is the one measure that is found to be unaffected by complexity.

Complexity induced a substantial reduction in trade efficiency. In the simple treatment, the average trade efficiency is 81%, which is consistent with the efficiency level found in previous experiments that use this bargaining mechanism (e.g., Schotter (1990)). In contrast, the efficiency in the complex treatment is only 73%. This difference is both economically and statistically significant. It is important to note that prior literature finds the level of efficiency to be a robust feature of the bargaining *institution* that we use in this study, not the environment. Schotter (1990), on page 222, states that: "These results substantiate the claim that the mechanism appears to be robust not only to the parameters of the environment but also to the manner in which people behave under it given these parameters." In this context, the reduction in efficiency we observe under complexity is meaningful.

One may speculate that the lower level of efficiency observed in the complex treatment could be attenuated by providing subjects more trading opportunities. That is, if subjects were given more chances to trade, perhaps they would adapt and efficient trading would arise. We find this not to be the case. We compare the difference in efficiency across periods with one and three rounds of bargaining and find that the treatment effect was significantly more pronounced when subjects received more trading opportunities (three-round periods). Table III reports the average efficiency level in one- and three-round periods in both treatments across the five sessions. The improvement in efficiency in the complex condition is neither economically (72% vs. 73%) nor statistically significant (p-value of 0.85 when conducting a mean average test treating each session as an observation). However, the improvement in efficiency in the simple condition is very pronounced—84% in three-round periods versus only 73% in one-round periods—and statistically significant (p-value of 0.03). Thus, more trading opportunities improved trade efficiency in the simple condition, but fail to do so when complexity is added. This evidence works against the intuition that more trading opportunities attenuate the effect complexity has on efficiency.

Table III
Experiment 1: Efficiency and Bidding Rounds

The table reports the average efficiency (subjects' payoffs from trade divided by available surplus) across sessions. We divide periods by number of trading rounds with one possible trading round in columns 1 and 2, and with three possible trading rounds in columns 3 and 4, and by complexity

	One-Round Periods		Three-Round Periods	
Session	Simple	Complex	Simple	Complex
1	0.89	0.80	0.88	0.80
2	0.51	0.86	0.78	0.44
3	0.63	0.64	0.84	0.76
4	0.75	0.72	0.94	0.73
5	0.70	0.51	0.82	0.94
Total	0.73	0.72	0.84	0.73

C. Bidding Behavior

To better understand what is driving our results, we study subjects' bidding strategies by estimating how their bids are driven by their estimates of their private values. We focus on linear bidding strategies since they are common in the experimental literature, and Radner and Schotter (1989) and Schotter (1990) suggest that subjects' bids are approximately linear in their private value when a linear equilibrium exists. Admittedly, equilibrium bidding behavior in the environment we study is more complicated than that in the independent private values setting studied in most of the experimental literature. However, this approach provides a simple measure of how traders respond to their information without imposing a likely unjustified structure on trader behavior. Notwithstanding, we show shortly that a linear estimation model has high explanatory power in this setting.

For each subset of observations, we estimated the following regression:

$$Bid_{i,t} = \beta_0 + \beta_1 \times Guess_{i,t} + \beta_2 \times Complex_t + \beta_3 \times Guess_{i,t} \times Complex_t + \epsilon_{i,t}, \quad (1)$$

where $Guess_{i,t}$ is subject i's guess of the security value in period t, and $Complex_t$ is a dummy variable representing the treatment in period t such that it is equal to one if the condition is complex and zero otherwise.

Table IV presents the estimation results, using robust regressions to control for outliers, and Figure 3 depicts the bid functions of buyers and sellers in the simple and complex conditions in one-round periods.[10] The first two columns in Table IV use data from periods in which there was only one bargaining round, and the second two columns use data from the first round in periods where three bargaining rounds were available. We find that the linear bidding function describes the data very well—the coefficient on subjects' guesses is

[10] If we handle outliers by winsorizing the top 5% of bids and use ordinary least squares (OLS) regressions, we obtain similar results to those obtained using robust regressions with no winsorization.

Table IV

Experiment 1: Bid Functions

We regress bids on estimated private values (guesses) and a treatment dummy using robust regressions. We estimate the model separately for buyers, sellers, one-negotiation-round periods, and three-negotiation-round periods. The first four columns include bids from all subjects. The last two columns include bids from subjects whose private value estimation error was similar across treatments. Standard errors are in brackets. Coefficients marked with ***, **, and * are statistically significant at the 1%, 5%, and 10% levels, respectively.

	One-Round Periods		Three-Round Periods		One Round—Constant Error	
	Seller	Buyer	Seller	Buyer	Seller	Buyer
Complex	4.009	4.762	5.965	4.663	4.042	8.743
	[1.741]**	[1.670]***	[1.328]***	[1.592]***	[3.110]	[2.104]***
Guess	0.738	0.691	0.574	0.609	0.950	0.649
	[0.108]***	[0.059]***	[0.074]***	[0.056]***	[0.191]***	[0.071]***
Complex × Guess	−0.439	−0.209	−0.461	−0.181	−0.413	−0.392
	[0.121]***	[0.065]***	[0.088]***	[0.069]***	[0.232]*	[0.084]***
Constant	8.006	2.612	9.788	2.805	6.680	3.240
	[1.511]***	[1.476]*	[0.978]***	[1.263]**	[2.628]**	[1.781]*
Observations	223	222	285	284	124	138
R^2	0.27	0.69	0.20	0.46	0.27	0.47

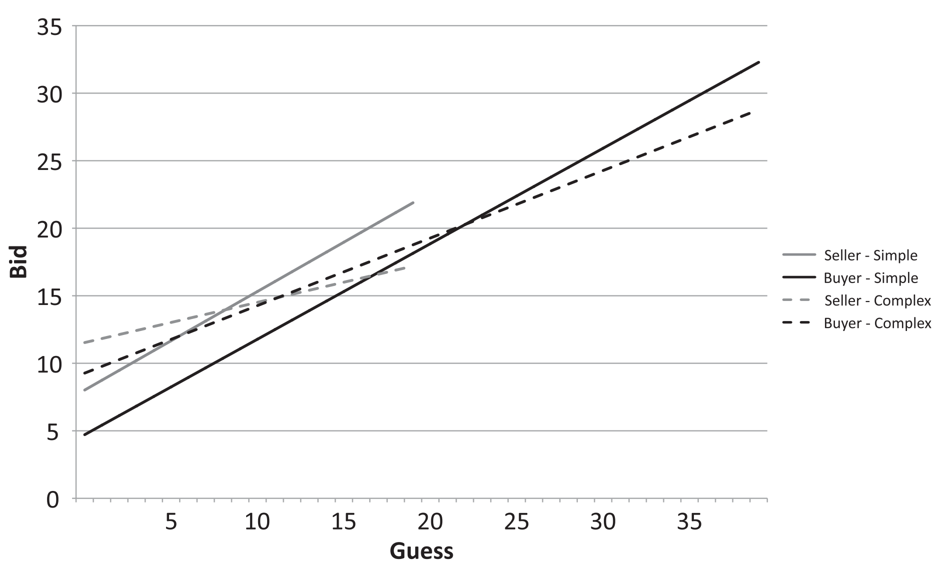

Figure 3. Experiment 1: Estimated bid functions (one-round periods). The figure depicts the estimated bid behavior as a function of private values. We plot the functions separately for buyers and sellers under simple and complex conditions, all for bidders participating in one-round bargaining games.

positive and statistically different from zero across all columns. In addition, the explanatory power of the model is quite high, with R^2's ranging from around 20% to 69%.[11] In addition, we find a strong treatment effect. In all four columns, the bid function is *less* responsive to the subjects' own guess of the value in the complex condition, compared with the simple condition.

To explore this further, we analyze a subset of our subjects who had a constant rate of estimation error across treatments. Our purpose here is to analyze the effect of complexity on bidding when estimation error does not vary. The results are in the last two columns of Table IV. Clearly, complexity did cause the bid function to be flatter, even when estimation error did not vary across treatments. This result further supports our previous claim that subjects' change in behavior in the complex environment is distinct from uncertainty about private values.

As intuition would suggest, we find that subjects' bid functions were less aggressive when they had more negotiation rounds. According to Table IV, the seller's private value guess coefficient in the three-round periods is 0.57, down from 0.74 in the one-round periods. Likewise, buyers' private value guess coefficient in the three-round periods is 0.61, down from 0.69 in the one-round periods. At the same time, the effect of complexity on the bid function appears to be similar—the coefficient interacting the complexity dummy variable with guesses is −0.44 for sellers in one-round periods and −0.46 in three-round periods. Likewise, the coefficient interacting the complexity dummy with guesses for sellers is −0.21 in one-round periods and −0.18 in three-round periods.[12]

To formally test whether bid functions are flatter in the initial round when subjects have three rounds of trade, we estimate the following regression for buyers and sellers separately: $Bid_{i,t} = \alpha + \beta_1 * ThreeRounds_t + \beta_2 * Guess_{i,t} + \beta_3 * Complex_t + \beta_4 * ThreeRounds_t * Guess_{i,t} + \beta_5 * Complex_t * Guess_{i,t} + \epsilon_{i,t}$. The results confirm the patterns observed in Table IV. The intercept of the bid function is higher for sellers in three-round periods compared with one-round periods (2.9, significant at the 5% level), while not for buyers. The coefficient is lower for both buyers and sellers in three-round periods (−0.15 and −0.09, respectively, both significant at the 5% level). Also consistent with the results reported above, the R^2's in these regression are higher for buyers than sellers (63% and 28%, respectively).

According to the first four columns in Table IV, the first round of bidding in these periods was characterized by less aggressive bidding by both buyers and sellers as compared with the case when only one round of bidding was available. One may expect subjects to improve their bids after each round of failed bargaining. To test this possibility, we create two additional measures:

[11] The difference in R^2's is consistent with the asymmetry in the strategic uncertainty faced by buyers and sellers. Buyers face less strategic uncertainty regarding the sellers' private values than the sellers do. As such, it is not surprising that buyers' bid functions are better correlated with their private values.

[12] These results are robust to the inclusion of subject fixed effects or estimation errors.

Table V

Bid Change across Rounds

We regress the bid change across rounds of negotiation (columns 1 and 2) and a dummy for bid improvement (columns 3 and 4) on the treatment dummy, number of periods, and their interaction. *Bid Change* equals round t minus round $t-1$ bid for buyers and round $t-1$ minus round t bid for sellers. *Bid Improvement* equals one if the bid change is weakly positive, and zero otherwise. In columns 1 and 2 we use OLS regressions and in columns 3 and 4 we use probit regressions. All standard errors (in brackets) are clustered by subject. Coefficients marked with *** and ** are statistically significant at the 1% and 5% levels, respectively.

	Bid Change		*Bid Improvement* (Weakly)		*Bid Improvement* (Strongly)	
Complex	0.414	−1.671	−0.010	−0.775	0.095	−0.554
	[0.423]	[0.710]**	[0.136]	[0.316]**	[0.134]	[0.282]**
Period		−0.048		−0.029		−0.012
		[0.049]		[0.024]		[0.022]
Complex × Period		0.288		0.106		0.095
		[0.092]***		[0.038]***		[0.035]***
Constant	1.504	1.950	0.634	0.913	0.351	0.459
	[0.345]***	[0.556]***	[0.125]***	[0.256]***	[0.120]***	[0.223]**
Observations	514	514	514	514	514	514
R^2	0.002	0.034	0.000	0.027	0.001	0.028

(i) *Bid Change:* current minus previous bid for sellers, and previous minus current bid for buyers. This quantity is positive when buyers and sellers improve their bids' competitiveness.

(ii) *Bid Improvement:* a dummy variable taking the value of one if the current bid is more competitive than the prior bid and zero otherwise. We consider both improvement in the weak sense (equal or greater) or in the strong sense (strictly greater).

Table V summarizes our results regarding the change in bids' competitiveness across rounds of bargaining. In the first two columns the dependent variable is *Bid Change*, and we use standard regressions. In the following four columns, the dependent variable is *Bid Improvement* and we use probit regressions altering the definition of improvement—in the weak sense (columns 3 and 4) and the strong sense (columns 5 and 6).

The results suggest that subjects indeed improved their bids in at least 63% of rounds, and that the average improvement size was 1.5 ECUs. Without controlling for differential change in behavior over periods for the two treatment conditions, we find that the cross-round improvement for subjects in the simple and complex treatments is similar. However, when we interact the period number with the treatment dummy, we find that initially subjects in the simple treatment were much more prone to increase their bids' competitiveness than subjects in the complex treatment (treatment effect of −1.67 on the magnitude of bid chance and −0.78 on bid improvement, defined weakly). At the same time, there is some evidence that the competitiveness of the bids in the complex treatment improved over time as the coefficient on the interaction between the treatment dummy and the period number is positive and statistically different from zero in all regressions. Finally, the results are qualitatively similar when we define bid improvement in the weak or strong sense. However, it is clear that in a large number of rounds subjects chose to keep their bids constant. For example, in the simple condition the frequency of bid improvements was 91% when defined weakly and only 46% when defined strongly.

These results are consistent with the efficiency findings in Table III. Table III suggests that the improvement in efficiency when more trading rounds are permitted is pronounced in the simple condition but not in the complex condition. Consistent with this result, we find in Table V that the bid improvement across rounds is much stronger in the simple condition compared with the complex condition. That is, in the simple condition, subjects are able to improve their bids in multiround periods resulting in more frequent transactions. However, we do not find similar evidence in the complex condition.

D. Adverse Selection

As the results in Section III.C indicate, both buyers and sellers shade their bids in the complex treatment compared to the simple treatment, even among subjects who make similar errors in the two treatments. This suggests that subjects' change in behavior in the complex environment is distinct from uncertainty about private values. On the other hand, behavior in the uncertainty treatment does not change between more and less uncertainty in a way that leads to reduced trading efficiency.

What, then, explains why noise does not cause changes in trading behavior, but complexity does? We believe that a likely source of trade breakdown comes from traders' estimations (or possibly misperceptions) of their own ability. Healy and Moore (2008) demonstrate theoretically and experimentally that agents tend to overestimate their relative ability when facing easy tasks and underestimate their relative ability when facing difficult tasks. They argue that, when faced with a hard task, agents are not able to determine if their difficulty in completing the task arises from the characteristics of the task or from their own lack of ability. A Bayesian agent facing a problem that is, in fact, more difficult than expected will lose confidence in his own ability relative to that of other agents.

In our setting, a complex problem might lead traders to question their own ability relative to others. Since private values are affiliated, buyers who cannot determine their own private value with confidence might fear that trade will only occur when the seller, who is presumed to be better informed, has a lower private value. Affiliation then implies that the uninformed buyer is likely to lose by trading against the informed seller. Note that this will hold even when the hypothetical informed seller believes that he trades against a fully informed buyer. That is, a buyer who loses confidence in his own ability will be concerned about trading against a skilled seller, even if the skilled seller believes himself to be trading against a skilled buyer.

When the problem facing both traders is in fact very difficult, as holds for our complex treatment, both players may mistakenly, but rationally, think that they have below average skill at the task in question and trade defensively. If this is the case, the buyer will shade his bid down, and the seller will shade his ask up. This produces something akin to two-sided adverse selection, and when the problem is sufficiently difficult relative to the expectations of the traders such concerns may override the potential gains to trade and lead to a breakdown in trade. This relates to the general idea that disparity in financial expertise can be destabilizing because it worsens adverse selection problems in trading venues (e.g., Glode, Green, and Lowery (2012)). As complexity increases, the probability that one opponent has an informational advantage over another rises. The two-sided adverse selection generated by uncertainty over the difficulty of the problem exacerbates this problem and can generate trade breakdowns even when both parties are, in fact, identically well informed.

One extreme response traders could employ in response to these concerns, and consistent with the theory, would be to refrain from trading altogether. To look at this in our data, we identify bids that reflect reluctance to trade. Specifically, we identify buyers' and sellers' bids in which there was no chance they could take a loss in a trade. These are buyers' bids that were at the lowest range of possible values (bids equal to or lower than one) and sellers' bids that were at the highest range of possible values (bids equal to or greater than 40). While these extreme bids were very infrequent in the data, we find that they were almost four times more prevalent in the complex condition than in the simple condition (1.52% of bids in the complex condition vs. 0.40% of bids in the simple condition). Comparing the subjects' propensity to submit no-trade bids across treatments, we find that the difference between complex and simple conditions was statistically different ($p = 0.083$). This evidence is consistent with the idea that under complexity subjects are more concerned about adverse selection and therefore are more reluctant to trade.

Thus, we find theoretical and empirical support for our proposed channel, which is consistent with our results on the effect of complexity on trade efficiency. The fact that, in contrast, simply adding noise to private values does not account for our findings is not surprising. Consider an extreme case in which a portfolio problem is so difficult that neither the buyer nor the seller can solve it. When neither player believes she has successfully solved the problem and neither player believes that the other player has successfully solved the problem, the auction is equivalent to one in which private values are common knowledge. This game, in contrast to the game in which players are assumed to correctly learn their individual private value, has fully efficient equilibria. Therefore, noise does not have to adversely affect trading efficiency and may in some instances improve it. In our setting, it appears to have an equivocal effect.

Table VI
Experiment 1: Learning

The table reports the average level of a number of dependent variables across period blocks (1–5, 6–10, 11–15) and treatment conditions. The dependent variables include the fraction of periods in which subjects provided an exact guess of their private value, the average guess error, the average payoffs, and the fraction of periods resulting in a transaction.

	Periods 1–5	Periods 6–10	Periods 11–15	Total
Simple Condition				
Frequency of exact guesses	0.72	0.68	0.76	0.72
Average guess error	2.85	3.02	1.42	2.38
Frequency of bid error	0.01	0.02	0.01	0.01
Average payoffs	3.77	4.51	3.64	3.98
Frequency of transaction	0.67	0.78	0.61	0.69
Complex Condition				
Frequency of exact guesses	0.32	0.29	0.33	0.31
Average guess error	9.48	5.58	7.27	7.55
Frequency of bid error	0.03	0.01	0.04	0.03
Average payoffs	3.67	3.33	3.72	3.58
Frequency of transaction	0.66	0.56	0.75	0.65

E. Robustness

Two possible concerns with the external validity of our experimental results relate to learning during the experiment or tiring out. One may worry that the results obtained in the laboratory reflect lack of experience with the task and that, upon repetition, subjects substantially alter their behavior. Alternatively, it might be possible that subjects become apathetic or bored during the experiment, making more errors as the 15 periods elapsed.

These are valid concerns and thus we consider whether learning or tiring out during the course of the experiment appeared to be significant factors. To that end, we analyze the variation of key measures across different stages of the experiment. Since the treatment condition varied across periods, we separate the data by condition. Table VI presents the results. Looking at various error measures, subjects' payoffs and the frequency of transactions do not vary monotonically with period number during the experiment. The only measure that appears to be significantly higher during the first set of periods (1–5) is the average guess error in the complex treatment. Therefore, while subjects might learn or alternatively tire out during the experiment, such effects do not appear to be a first-order concern.

Another possible concern might be that variation in attendance in different sessions might affect our results. Recall that our within-subject design is predicated on observing a difference in behavior between the same subjects across differences in security complexity. To avoid possible confusion between the main treatment and learning, we vary the order of treatments within a session. Given that the number of participants varied across session based on attendance, the number of observations collected from simple periods in the first block (periods 1–5) is not the same as the number of observations collected from complex periods in the same block. To control for this difference, we conduct a robustness check in which we repeat the analysis but randomly remove observations such that within a block of periods (1–5, 6–10, 11–15), the number of simple and complex periods matched.

The results of this robustness check are in the Internet Appendix. Briefly, we find very similar results. The dependent variables (frequency of guess errors, magnitude of guess errors, and frequency of bid errors) are very similar in magnitude to the main results. Likewise, the difference across treatment cells yields the same statistical significance. The same is generally true for our measures of efficiency, liquidity, and price volatility with some minor differences. The differences in efficiency and price volatility across treatments appear to be more pronounced across treatment when equalizing the number of observations across cells, while the difference in bid gap appears to be less pronounced. The reduction in the number of observations naturally leads to more variation in levels, resulting in somewhat weaker p-values across some of the measures. This holds true for the second experiment as well.

To further assess the robustness of our results with respect to the unbalanced number of observations, we also reevaluate the effect of this on our estimates of the bidding strategies. Specifically, we take the same data set described above, in which the number of observations across blocks/treatments is matched, and estimate subjects' bidding strategies as in Section III.C. As the results in the Internet Appendix show, across the four regressions (one-round seller, one-round buyer, three-round seller, three-round buyer) the results obtained from this smaller data set are very similar to those obtained with the full data set. That is, the bidding function intercept and slope are very similar, and, more importantly, the previously documented effect of complexity on the bid functions is very similar.

IV. Conclusions and Discussion

Complexity of financial markets has become a fact of life. Based on recent events during the 2008 financial crisis, it is now clear that financial models can fail and that market participants are not all-knowing. Previous crashes and crises have also made such points, but our profession appears to be more receptive to investigating the effects of complexity on markets at the present time.

In this paper, we study the effects of complexity on trading behavior in an experimental market setting. We show that complexity leads to lower liquidity and efficiency, and higher price volatility. Strikingly, this appears to be separate from the added noise that estimation errors impose on asset trading.

We view our work in this paper as a first step to understanding how complexity affects markets and informs policy. Indeed, our results tend to support the policy implication that standardization of assets may improve welfare, making assets more liquid and less volatile. Admittedly, whether traders in real markets have the same response to complexity remains an open, but important, question. In addition, understanding the effect of complexity on behavioral biases warrants future investigation. Evaluation of this effect, as well as its theoretical underpinnings, is the subject of future research.

Initial submission: January 14, 2013; Final version received: January 22, 2013
Editor: Campbell Harvey

REFERENCES

Arora, Sanjeev, Boaz Barak, Markus Brunnermeier, and Rong Ge, 2010, Computational complexity and information asymmetry in financial products, Working paper, Princeton University.

Bernardo, Antonio E., and Bradford Cornell, 1997, The valuation of complex derivatives by major investment firms: Empirical evidence, *Journal of Finance* 52, 785–798.

Brunnermeier, Markus, and Martin Oehmke, 2010, Complexity in financial markets, Working paper, Princeton University.

Carlin, Bruce I., 2009, Strategic price complexity in retail financial markets, *Journal of Financial Economics* 91, 278–287.

Carlin, Bruce I., and Gustavo Manso, 2011, Obfuscation, learning, and the evolution of investor sophistication, *Review of Financial Studies* 24, 754–785.

Chatterjee, Kaylan, and William Samuelson, 1983, Bargaining under incomplete information, *Operations Research* 31, 835–851.

Fischbacher, Urs, 2007, z-tree: Zurich toolbox for ready-made economic experiments, *Experimental Economics* 10, 171–178.

Furfine, Craig, 2011, Deal complexity, loan performance, and the pricing of commercial mortgage backed securities, Working paper, Northwestern University.

Glode, Vincent, Richard Green, and Richard Lowery, 2012, Financial expertise as an arms race, *Journal of Finance* 67, 1723–1760.

Healy, Paul J., and Don A. Moore, 2008, The trouble with overconfidence, *Psychological Review* 115, 502–517.

Healy, Paul J., John O. Ledyard, Sera Linardi, and J. Richard Lowery, 2010, Prediction markets: Alternative mechanisms for complex environments with few traders, *Management Science* 56, 1977–1996.

Kadan, Ohad, 2007, Equilibrium in the two-player, k-double auction with affiliated private values, *Journal of Economic Theory* 135, 495–513.

Payzan-LeNestour, Elise, and Peter Bossaerts, 2011, Learning to choose the right investment in an unstable world: Experimental evidence, Working paper, California Institute of Technology.

Radner, Roy, and Andrew Schotter, 1989, The sealed-bid mechanism: An experimental study, *Journal of Economic Theory* 48, 179–220.

Schotter, Andrew, 1990, Bad and good news about the sealed-bid mechanism: Some experimental results, *American Economic Review* 80, 220–226.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's web site:

Appendix S1: Internet Appendix