

Strategic Bidding of Offer Curves:
An Agent-Based Approach to Exploring Supply Curve Equilibria

Steven O. Kimbrough
University of Pennsylvania

Frederic H. Murphy
Temple University

March 11, 2011

Abstract

We model a market in which suppliers bid step-function offer curves using agent-based modeling. Our model is an abstraction of electricity markets where step-function offer curves are given to an independent system operator that manages the auctions in electricity markets. Positing an elementary and computationally accessible learning model, `PROBE AND ADJUST`, we present analytic results that characterize both the behavior of the learning model and the properties of step-function equilibria. Thus, we have developed a framework for validating agent-based models prior to using them in situations that are too complicated to be analyzed using traditional economic theory. In addition, we demonstrate computationally that, by using alternative policies, even simple agents can achieve monopoly rewards for themselves by pursuing more industry-oriented strategies. This raises the issue of how participants in oligopolistic markets actually behave.

Contents

1	Introduction	1
2	Supply Function Equilibria	2
3	Agent-based Models	4
4	Representation of the Market and the Agents	6
5	Agent Behaviors	8
5.1	Notation	8
5.2	Iterations with Agents on the Verticals of the Offer Curve (Quantities)	10
5.3	Iterations when One Player Sets the Price to Block a Supply Step of Another Player	12
5.4	Iterations with One Player setting the Trial Market Price by its Offer Price and the Other on the Vertical of Its Offer Curve	13
6	Existence and Uniqueness	15
7	Exploring the Analytic Results with the Computational Model	18
7.1	On the Verticals	19
7.2	Bertrand Competition	19
7.3	Mixed-Condition Equilibrium	20
8	Further Explorations with the Computational Model	22
8.1	P2Both2Is25 with 6 Bid Steps	22
8.2	Q2Firm1Is250	24
8.3	Bertrand250Both 6 Steps	24
8.4	Multiple Equilibria	26
9	Exploring a Non-Myopic Strategy: MR-COR	26
9.1	Beyond on the Verticals	26
9.2	Beyond Bertrand	27
9.3	Beyond the Mixed-condition Equilibrium	28
9.4	Q2Firm1Is250	30
9.5	NFirmsSimple	30
10	Summary and Conclusion	30
	References	32

A Probe and Adjust in Pseudo-Code	36
B Proofs	37
B.1 Lemma 1	37
B.2 Proof of Lemma 2	37
B.3 Proof of Theorem 1	38
B.4 Proof of Corollary 1	38
B.5 Proof of Lemma 3	39
B.6 Proof of Lemma 4	39
B.7 Proof of Theorem 2	39
B.8 Proof of Lemma 5	39
B.9 Proof of Lemma 6	40
B.10 Proof of Lemma 7	40
B.11 Proof of Lemma 8	41
B.12 Proof of Lemma 9	41
B.13 Proof of Theorem 3	42
B.14 Proof of Lemma 10	42
C Computational Model	43
D Additional Computational Results	47
D.1 Run Summaries	47
D.2 Screen Dumps	48

List of Figures

1	Player 1 definitions.	10
2	An equilibrium where the first steps of both players are utilized and player 1 raises its offer price to satisfy its Cournot equilibrium condition and the price P_1^2 is below p_2^2	15
3	“Supply & Demand Curves” chart with the Base Case Parameterization. From Figure 20 in the supplemental materials	18
4	P2BothIs25 Parameterization, supply and demand plot for a representative run (both firms update with Own Returns), with both firms starting with high bids. From Figure 29 in §D.	19
5	Bertrand250Both Parameterization, prices for a representative run (both suppliers update with Own Returns). From Figure 21 in §D.	20
6	PSBothIs25 Parameterization, profits for a representative run (both firms update with Own Returns). See Figure 24 in §D for another example.	21
7	PSBothIs25 Parameterization, supply and demand plot for a representative run (both firms update with Own Returns). See Figure 24 in §D for another example.	21
8	P2BothIs25 parameterization, but with 6 bid steps for each firm, prices for a representative run (both suppliers update with Own Returns), each firm initializing with low bids. From Figure 28 in §D.	23
9	P2BothIs25 parameterization, but with 6 bid steps for each firm, profits for a representative run (both suppliers update with Own Returns), each firm initializing with low bids. From Figure 28 in §D.	23
10	P2BothIs25 parameterization, but with 6 bid steps for each firm, supply and demand curves for a representative run (both suppliers update with Own Returns). From Figure 28 in §D.	24
11	P2BothIs25 parameterization, but with 6 bid steps for each firm, supply and demand curves for a representative run (both suppliers update with Own Returns), firm1 initialized to 16 for prices, firm 2 to 100. From Figure 33 in §D.	25
12	Bertrand250Both 6 Steps setup, representative run. Each firm has 6 bid steps. All firms start by bidding low. From Figure 35 in §D.	25
13	Bertrand250Both Parameterization, prices for a representative run (both suppliers update with MR-COR). From Figure 22 in §D.	27
14	Bertrand250Both Parameterization, supply and demand curve for a representative run (both suppliers update with MR-COR). From Figure 22 in §D.	27
15	Bertrand250Both Parameterization, profits for a representative run (firm 1 updates with Own Returns, firm 2 with MR-COR). From Figure 23 in §D.	28
16	P2BothIs25 Parameterization, profits plot for a representative run (both firms update with MR-COR). From Figure 25 in §D.	29

17	P2BothIs25 Parameterization, supply and demand plot for a representative run (both firms update with MR-COR). From Figure 25 in §D. See also Tables 9 and 10 in the supplemental materials.	29
18	NFirmsSimple parameterization, plots for a representative run. Five firms update with MR-COR; each firm has 6 bid steps and two plants of capacity 15 and 55. All firms start by bidding low. QIntercept is 300. From Figure 34 in §D.	30
19	Pseudo-code for basic PROBE AND ADJUST	36
20	Base Case Parameterization, representative run (both suppliers update with Own Returns) .	45
21	Bertrand250Both stored case, representative run (both suppliers update with Own Returns) .	51
22	Bertrand250Both stored case, representative run (both suppliers update with MR-COR) . . .	52
23	Bertrand250Both stored case, representative run (firm 1 updates with Own Returns, firm 2 with MR-COR)	53
24	P2BothIs25 stored case, representative run (both suppliers update with Own Returns)	54
25	P2BothIs25 stored case, representative run (both suppliers update with MR-COR)	55
26	P2BothIs25 stored case, representative run (firm1 updates with Own Returns, firm 2 with MR-COR)	56
27	Base Case parameterization, but with low starting prices for all firms and segments.	57
28	P2BothIs25 6 steps stored case, which is P2BothIs25, but with 6 bid steps for each firm . . .	58
29	P2BothIs25 stored case, but both firms start bidding high	59
30	P2BothIs25 6 steps stored case, which is P2BothIs25 setup, but with 6 steps each firm, both firms start bidding high at 100 each step	60
31	Q2Firm1Is250 6 steps stored case, which is the Q2Firm1Is250 setup, but with 6 steps for each firm, both firms start bidding low at 16 each step, a representative run	61
32	Q2Firm1Is250 stored case, representative run (both firms update with Own Returns)	62
33	Q2Firm1Is250 setup, representative run (both firms update with Own Returns), but both firms have 6 bid steps. Firm 1 starts at base prices = 16, firm 2 at 100.	63
34	NFirmsSimple setup, representative run. Five firms update with MR-COR; each firm has 6 bid steps and two plants of capacity 15 and 55. All firms start by bidding low. QIntercept is 300.	64
35	Bertrand250Both 6 Steps setup, representative run. Each firm has 6 bid steps. All firms start by bidding low.	65
36	Q2Firm1Is250 Parameterization, representative run (both suppliers update with MR-COR) .	66
37	Q2Firm1Is250 Parameterization, representative run (S1 updates with Own Returns, S2 with MR-COR)	67

List of Tables

1	Default settings of the environment parameters, Base Case Parameterization.	43
2	Default settings of the agent parameters, Base Case Parameterization.	44
3	Base Case Parameterization (both suppliers update with Own Returns). See Figure 20 in §C.	48
4	Bertrand250Both Parameterization (both suppliers update with Own Returns). See Figure 5.	48
5	Q2Firm1Is250 Parameterization (both suppliers update with Own Returns). See Figure 32 in §D.	48
6	Base Case Parameterization (both suppliers update with MR-COR).	49
7	Base Case Parameterization (firm 1 updates with Own Returns, firm 2 with MR-COR). . . .	49
8	P2BothIs25 Parameterization (firm 1 updates with Own Returns, firm 2 with MR-COR). See Figure 26 for a representative run.	50
9	Q2Firm1Is250 Parameterization (both suppliers update with MR-COR). See Figure 36 in §D.	50
10	Q2Firm1Is250 Parameterization (firm 1 updates with Own Returns, firm 2 with MR-COR). See Figure 37 in §D.	50

1 Introduction

In this paper we characterize the pure-strategy supply function equilibria when the players can set both prices and quantities on step-function offer curves. We present an agent-based approach that models possible outcomes resulting from agents making step-function offers with the goal of understanding market outcomes and connecting classical economic theory with agent-based modeling. We develop our model in the context of oligopolistic markets, using both analytic results and simulations. While the offer curves in markets such as electricity auctions are step functions, the literature on supply-function equilibria has mostly focused on continuous supply functions, creating a disconnect between the theoretical literature and the actual bidding behavior of suppliers. Connecting agent behaviors and simulated market outcomes to economic theory is essential to avoiding ad hoc modeling leading to misleading policy conclusions.

The two traditional models of oligopoly are the Bertrand model, in which firms bid prices and the firm with the lowest price wins all the demand at that price, and the Nash-Cournot model, in which firms bid quantities and the market price is set so that the market clears at the sum of the quantities offered. On the surface, an electricity auction seems to be modeled best as a Bertrand market because bidders offer prices. However, capacity is limited and as Kreps and Scheinkman (1983) have shown, in two-stage games where capacity is determined in a first-stage game followed by bidding prices, Nash-Cournot and Bertrand lead to the same equilibrium. Fabra, von der Fehr, and Harbord (2006) show ways in which capacity limits and step functions alter the Bertrand equilibrium in a duopoly. For uniform-price auctions they conclude that the equilibrium price is at the higher of the two player's costs when demand is low and at the price cap when demand is high.

If the Nash-Cournot and Bertrand equilibria were to lead to the same equilibrium, one would presume that a step function should produce the Nash-Cournot equilibrium. As a counter to this result, because of the load curve and uncertainty in demand, there are times when capacity is not binding for at least one of the players. Edgeworth (1925) shows that in this case there is no stable equilibrium and the market outcome is a cycle. Kruse et al. (1994) use laboratory experiments with human subjects and find that a model with adjustments in the direction predicted by Edgeworth is probably the best explanation of the experimental outcomes. Ernst, Minoia, and Ilic (2004) use agent-based models to model capacity-limited markets and verify Edgeworth's result.

As an emerging modeling approach, agent-based models have the promise of offering insight into economic and social outcomes in situations that are too complex to be represented by the tractable models required to prove theoretical results. However, for agent-based approaches to attain legitimacy in the larger modeling community, the properties of the agent behaviors need to be understood. The results of simulations should replicate the theoretical models when using the same assumptions as the theory and the theoretical properties of the agent's actions need to be developed.

In this paper we present an agent-based model for a market in which supplier agents bid step-function offer curves and we explore the search outcomes of the agents in their quests to improve their profits.

Central to our model of agents is their learning regime, which is a local search heuristic we term PROBE AND ADJUST. We prove that these agents move towards equilibria that satisfy Cournot and Bertrand equilibrium conditions, when these equilibria are applicable. There is also a third possible equilibrium that the agents find where one player satisfies the Cournot equilibrium condition and the other sees marginal revenue equal to price and winds up at one of the quantity break points in its marginal-cost step function. The convergence results for PROBE AND ADJUST serve a dual role in characterizing the possible pure-strategy equilibria. We then show that an equilibrium always exists and there can be multiple equilibria.

Thus, this paper addresses three questions. First, can agents using PROBE AND ADJUST replicate classical economic theory? Second, can these agents offer insight into the behavior of supply function equilibria and market power? Third, can the behavior of agents be validated so that simulations using the agent are more than ad hoc exercises? We answer all questions in the affirmative. In the process of answering the first question, we also add to economic theory. Our approach to characterizing agent behavior provides a template for assessing the behaviors of agents in general against an existing body of theory. This opens the opportunity to gain deeper insights into the design issues for complicated markets such as electricity markets and have greater confidence in results from simulations.

2 Supply Function Equilibria

Klemperer and Meyer (1989) developed the supply function equilibrium model to deal with uncertainty in demand because different demand levels have different equilibrium prices and a single bid of just a price or just a quantity is not optimal. In the case of electricity, demand varies widely throughout the day. A standard representation of demand in a day, month or year is a load duration curve where hourly demands are sorted from largest to smallest. The load duration curve has the form of an inverse probability distribution function, making the results of Klemperer and Meyer relevant even in the deterministic case. Green and Newbery (1993) recognized that the approach of Klemperer and Meyer is a more realistic representation of electricity markets than either the Nash-Cournot or Bertrand models because of the load duration curve and the need to bid a supply function. See Holmberg and Newbery (2009) for an overview of supply function equilibria.

Unfortunately, the supply function equilibrium suffers from several problems. Foremost are that the model has multiple equilibria and the equilibria are difficult to characterize. That the equilibria are difficult to characterize is not surprising because a supply function has many of the features of representing conjectural variation with the conjectural variation being consistent with the equilibrium solution. See Murphy and Smeers (2005) for the kinds of problems created by representing the opponent's response across stages in a two-stage game. Having multiple equilibria creates other problems. When comparing two policy scenarios one has to be careful to compare "equivalent" equilibria across scenarios if a concept of equivalent can even be defined.

Klemperer and Meyer recognized that multiple equilibria exist. A stream of literature has tried to narrow the set of equilibria. Baldick, Grant, and Kahn (2004) provide a useful summary of much of this literature. Baldick and Hogan (2002) show that any unconstrained nonlinear concave or convex supply-function equilibrium is unstable when the underlying cost function is affine. That is, the only supply functions in a stable equilibrium are affine. Baldick and Hogan (2002) and Holmberg (2008) show that with a price cap the supply function equilibrium is unique. Baldick and Hogan (2004) show that polynomial approximations to the supply functions can be stable only for polynomials with low powers. Rudkevich (2003) shows that a tâtonnement process where players set quantities does not converge when players have full information. Holmberg, Newbery, and Ralph (2010) find the pure strategy equilibria when prices but not quantities are adjusted in step-function offer curves.

In simulations Rudkevich gets convergence to the linear supply function equilibrium of Green (1997) when players estimate the supply functions of other players. Hobbs et al. (2000, 2004a, 2004b) find solutions using conjectural variation in transmission pricing. This last set of papers illustrates a fairly common property of conjectural variation models: because the conjectural variation does not have to be consistent with the players' behavior, it is easier to get an equilibrium as opposed to when the perceptions of variation have to be consistent with the opponent's response. To further illustrate the limited nature of the results, despite the complexity of the proofs, virtually all of the results on supply function equilibria are for symmetric games where the players have the same cost curves. Fabra, von der Fehr, and Harbord (2006), Anderson and Hu (2008), and Holmberg (2009) are exceptions.

The shape of the supply function is an important consideration and the possibility of non-monotonic supply functions exists. Anderson and Philpott (2002) show that if the opponent's supply function is known and has a log concave inverse then there is an increasing supply function that gives an optimal response for the player for every realization of the demand shock. Anderson and Xu (2005a) extend these results to include contracts. Both of these papers are for symmetric equilibria. Anderson and Hu (2008) extend their results to the case where the cost functions are not symmetric.

Von der Fehr and Harbord (1993) raise the issue that the supply functions in electricity bids are step functions. They develop randomized strategies to find an equilibrium with step functions and state that there are no pure-strategy equilibria. Anderson and Xu (2005b) note that in Australia the maximum number of prices is set at 10. They develop ϵ -optimal supply functions in response to a discrete supply function of the other player. Because the other player uses a discrete supply function, the derived demand distribution is discontinuous. They developed the optimal response but not the equilibrium solution. Wilson (2008) characterizes the equilibrium with uncertain demand and transmission congestion. He points out that his results are not especially encouraging because the equilibrium is dependent on random shocks and the equations are highly nonlinear.

Modeling electricity markets has been one of the drivers of the research on supply function equilibria because that is the structure of the bidding in the auctions. For a survey of the modeling literature, including

some agent-based models, see Ventosa et al. (2005).

Brandts et al. (2008) use human subjects in a laboratory setting to examine the differences between outcomes where some players are asked to play a Cournot game and others can bid supply functions. They also include a set of experiments with players being allowed to sell in futures markets. They found that having players bid a separate price for each unit of production offered leads to greater total production and a futures market increases production as well.

3 Agent-based Models

An alternative to analytic results is to simulate the market. Rudkevich (2003) did this to explore the stability of the supply-function equilibria in the context of a regional electricity market. He used the classic Cournot assumption of the player presuming the other player does not respond and the player knows the demand curve. Contreras et al. (2002) use a cobweb bidding model. Visudhiphan (2003) uses multiple behavioral rules for agents in simulating the New England energy market using step-function offer curves. An attempt to calibrate the results with actual outcomes was not successful because of the complexity of the actual market. Weidlich (2008) was successful in building an agent-based model of German electricity markets and replicating actual outcomes. These approaches simulate agents who are optimizing the standard objective function of an oligopolist, that is, maximizing profits given the positions of the other player and full information about the market. When these models converge, the resulting solution is an equilibrium that satisfies the conditions for one of the classic economic models.

Agent-based models differ from the classic optimizations in oligopoly theory in that they focus on agents that can learn in limited ways. The agents are allowed to make decisions with less than optimal returns in any one period but can lead to learning that results in long-run increased levels of profitability. Agents in agent-based models, thus typically embody some form of so-called distributed rationality; they need not be myopic. One can think of agent-based models as partial tests of rational decision making in that a smart person should be able to do better than an intelligence-limited agent that looks to improve its situation without the same knowledge as humans. If an agent can deliver high profits, a person should be able to do at least as well. Yet, if an agent cannot deliver high profits, it may be too limited and a smart person could potentially do better. Another difference is that agents do not necessarily have full information. Instead of acting with complete input information, in our case they act on outcomes. This fits well within the context of games with repeated play that have players learning as they play.

One of the important features of markets such as electricity markets is that the auction takes place daily for each hour in the day. Thus, it is a game with repeated play (with fixed, rather than varying counterparties, a condition we shall assume unless otherwise noted). As Rothkopf (1999) notes, this presents an opportunity for learning and cooperating as in repeated-play prisoner's dilemma. Agent-based modeling captures learning through explicit learning models of what agents know and remember and how they process

the information.

Krishna and Ramesh (1998a and 1998b) built an agent-based model where avatars negotiate profitable coalitions. This study is problematic from an electricity perspective because members of coalitions negotiate market relationships. This would violate the market rules of the system operator. Lave et al. (2004) develop a model with agents using a genetic algorithm to optimize supplier outcomes in electricity markets and also run scenarios where customers are allowed to collude. They find that when the customers are not allowed to collude and the objective functions of the suppliers maximize profits, the result is the competitive equilibrium. When the suppliers have objective functions that include a weight on the price level, they get close to the monopoly solution. When customers are allowed to collude and have their own genetic algorithm, they can force the price back down. Bower and Bunn (2000) were the first to apply agent-based modeling to electricity markets in Britain. Bunn and Oliveira (2001, 2003, 2008) expand the model and add a representation of the revised market rules under the New Electricity Trading Arrangements.

In all of the models mentioned here (in distinction to our model), players bid a single price and not a supply curve, and thus, are subject to Bertrand rules. Bertrand models can exhibit pricing above marginal cost only when a new feature is introduced into the market structure. For example, Boone (2002) lets suppliers make a sequence of price moves before the consumer sees the price. This model applies to airlines using their databases and a stream of prices to draw out competitor reactions to pricing moves before stabilizing on the prices almost all customers see. These moves take place at a time when few people are making reservations. Bagnall and Smith (2005) build on the work of Bunn and Oliveira with a different learning model and explore possible market outcomes. Because of capacity limits, Bagnall and Smith (2005) find patterns that are close to Edgeworth cycles and players have profits above marginal cost. The random moves of their players lead to solutions that do not replicate Edgeworth's results exactly.

Sun and Tesfatsion (2007) have built an agent-based model that incorporates a training model used by the New England and the PJM independent system operators (ISO's). The model has five generators and a grid that can become congested. Using reinforcement learning, they find that agents are able to raise their prices above marginal cost. The agents stabilize on an equilibrium and do not replicate the Edgeworth results because they change the upper and lower bounds on capacity offers.

The role of agent-based modeling in understanding economic phenomena can be understood by analogy to simulation and queuing theory. Queuing theory is useful for developing the theoretical properties of queues and first-cut results for a queuing system. However, real queuing systems are in general more complex than the mathematical models and the analytically tractable probability distributions often do not fit the data. Thus, real systems are studied mainly through the use of simulation models.

The same applies to markets. The theoretical models do not truly represent the institutions and uncertainties faced by the players. Players have opportunities for learning and the specifics of a market's microstructure can create opportunities to increase profits. We see this with supply-function equilibria. The theoretical results are for very stylized cases such as cost symmetry, continuous offer curves, and one-shot

games rather than games with repeated play. Thus, these results are not robust in the face of a very complicated market such as the one for electricity. Agent-based models can incorporate much higher degrees of market complexity. They hold great promise for being fully adequate for both theoretical and practical purposes. However, they have to be built carefully so that the agents' behaviors are either consistent with the assumptions of economic theory or are reasonable extrapolations of observable behavior. The same is true for their knowledge sets, what they know when making decisions.

One strategy is to have agents learn the derived demand curve and optimize based on that information. This is what Young (2009) observes in simulations using Cal Tech students as players. These assumptions lead to the classical results of oligopoly theory in both the Cournot and Bertrand games.

Newbery (2005) points out that Bertrand duopoly players can cooperate by alternating who has the lowest price, allowing players to capture economic rents. The players can also declare trigger strategies where they communicate their willingness to punish deviations from cooperation through a response of drastic price cuts for extended periods (Tirole, 1988). This makes the one-period gain from defection less than the long-run losses. Liski and Montero (2006) point out how futures markets can facilitate cooperation in a Bertrand game with repeated play through making trigger strategies (see Tirole, 1988) more damaging to the defector.

Kimbrough and Murphy (2009) show that players can use experimentation and learning to develop strategies in the classic Cournot game that lead either to splitting monopoly profits or achieving the Nash-Cournot equilibrium, depending on their choice of objective functions. The agents are endowed with minimal intelligence and do not communicate. So, the excess profits fit within the constraints of antitrust law.

We now extend our analysis to the bidding of offer curves. Each player bids multiple steps, with the number of steps capped a priori. This is analogous to the Australian and several U.S. electricity markets. In the bid the player specifies the quantity and price for each step.

Central to our results is establishing the properties of the PROBE AND ADJUST agent. Showing that this agent replicates economic theory in the standard models is a way to validate the use of this agent in economic situations that do not fit the standard model and the first step in using it to model more complicated markets. These results establish a template for one approach to validating the design and choice of an agent as part of building agent-based models. We also add to economic theory by showing that the classic results hold with agents that do not have full information and act myopically, unlike the classical assumptions. Then we show how the agent behaviors differ with alternative information and incentives.

4 Representation of the Market and the Agents

We have structured our model to come close to the day-ahead market in electricity auctions, a main area where supply-function equilibria are of interest. In the day-ahead market each firm bids a step-function offer curve, load-serving entities bid step-function demand curves, and the independent system operator solves a linear program with these supply and demand curves and transmission constraints to find the

market equilibrium. In our model we use a continuous demand curve and do not include the transmission constraints. Thus, the model can be solved by simply intersecting a single, aggregate offer curve with a single demand curve. The analysis is for a duopoly. However, we present simulations with more players.

Our agent implementation models three different objective functions for the firm, which allows us to compare simulations with the different objectives. The first is the classic in economics where a firm maximizes its own profits. In the second, the firm looks to maximize industry profitability, that is, the total profits of all players rather than just its own. Here the firm looks out for the best interest of its industry; this objective is principally interesting because of its role in the third objective. The third objective function is the same as the second except that the firm also requires that it maintain a “fair share” of the market. In essence, the companies using this rule have a convention on what a fair share is and apply a trigger strategy when they receive less than their fair share. In our case, a fair share is achieved when all firms have approximately the same fraction of the market. The last two objectives are consistent with the Kuenne’s (1999) view of firms forming a community with community norms about what is appropriate behavior.

The standard approach to modeling learning by the agents is to discretize the decision space into a finite set of choices, put a probability of selection on the choices, and use random numbers to select the decision. See Erev and Roth (1998). We use a different approach. We leave the decision space continuous, give each agent a baseline trial solution, and let the player randomly modify the baseline values by small amounts.

In each round of play, termed an *episode*, each player constructs an offer curve, the system (mimicking the ISO in an electricity market) sums the offer curves, producing an aggregate offer curve for the episode. The system then determines the price and quantity (based on the demand curve and the summed offer curves), and the market is cleared. Each player has a given marginal cost curve and constructs the actual offer curve by using randomization to vary the prices and quantities on the steps of the offer curve, taking care to keep the points on the offer curve at or above the underlying marginal cost curve. The randomizations of each player are independent of the randomizations of the other. After a fixed number of episodes, termed an *epoch*,¹ players assess the outcomes for all of their price and quantity increases and decreases. For each parameter (price or quantity), the direction of random variation that leads to the highest increase in the objective function determines the direction of change in the baseline offer curve for the next epoch. The players then adjust the prices and quantities by a small amount. We term our version of it PROBE AND ADJUST. Figure 19 in §A of the supplementary materials provides pseudo-code for quantity adjustment in the basic version. The adjustment for prices follows the same logic. In our implementation, players are deciding and learning both on prices and quantities, using PROBE AND ADJUST.

Note that we are using this algorithm in a strategic decision making context, rather than, as usual, in a parametric decision making context. Note also, this form of local search has features similar to those of optimization methods such as steepest ascent while the Erev and Roth (1998) approach falls into the category of global optimization methods for nonconvex functions. In the equilibrium context PROBE AND

¹Epoch lengths may vary among players and need not be synchronized.

ADJUST is similar to successive overrelaxation algorithms. Nevertheless, our main interest in PROBE AND ADJUST is that it is a plausible representation of the kinds of incremental actions taken by individuals and organizations.

5 Agent Behaviors

Here we show the different equilibrium points towards which the agents move when they have an objective of maximizing their own profits. In practice the agents can start on one path towards one equilibrium and through random outcomes switch to another path heading towards a different equilibrium. The limiting behavior described here is useful for characterizing how the agents move towards potential equilibria and what these potential equilibria are.

Our results on agent behavior can be considered characterizations of local equilibria in standard static analysis with much weaker assumptions on what the economic agents know. These results give us confidence that when we change some aspect of the agents' environment or objectives and they do not move towards the standard equilibria, the new equilibria are meaningful as alternatives to the classical results. Using agents allows us to incorporate into the models the complexities of institutions that cannot be characterized analytically while our analysis of agent properties with the standard assumptions of oligopoly theory gives us confidence that the outcomes with the added complexity are not artifacts of unexplored properties of the agents.

In a new episode, as the players vary the price and quantity breakpoints of the steps in the offer curves, the episode trial solution can have one of several properties. If these properties repeat, then the agents move to local equilibria with these properties.

If both players are at their quantity breakpoints in their offer curves, then the players see the breakpoint quantities, set the trial equilibrium and see themselves moving to a local Cournot solution. In the classic Bertrand game one player captures the market while the other is at the quantity breakpoint of zero in a one-step function. If both players offer large steps that can satisfy demand, then each player can set prices and quantities and the players move to the Bertrand equilibrium. We also reproduce the Edgeworth cycles observed by Ernst, Minoia, and Ilic (2004), though through a different mechanism. In other situations, we have one player setting the price and seeing the impact of price adjustments while the other sees the effect of its quantity adjustments. This leads the agents to move towards an equilibrium where the equilibrium condition for the price-setting player is the Cournot condition and the quantity-adjusting player sees marginal revenue equal to price. If both players offer the same price for a step that sets the equilibrium price, then the quantity split is indeterminate. The episodes of this type are unlikely to repeat.

5.1 Notation

We use an inverse demand curve $p = a - bQ$. Let

i index the players

$-i$ indicate the player that is not player i

k index the epochs

j index the episodes in an epoch

l index the steps on the offer and marginal cost curves

q_i^l be the quantities that define the discontinuities of the step function of the marginal cost curve of player i and are total supply through the l^{th} step.

c_i^l be the marginal cost on the cost-curve for breakpoint l

Q_i^k be the baseline quantity produced at the trial solution for epoch k by player i

P_i^k be the baseline offer price of player i at epoch k that is either setting the equilibrium price or is closest to setting the price.

ϵ be the adjustment for trial quantities at the end of each epoch

$b\epsilon$ be the adjustment for trial prices at the end of each epoch

δ be the amount by which quantities and prices randomly change up or down in each episode

J be the number of episodes in an epoch

The marginal cost is c_i^l for any total demand for i 's product q such that $q_i^l \leq q < q_i^{l+1}$. At episode k player i constructs an offer curve with step l defined by $(q_i^{k,l}, p_i^{k,l})$ with $q_i^{k,l} < q_i^{k,l+1}$ and $p_i^{k,l} < p_i^{k,l+1}$. The price is $p_i^{k,l}$ for any q such that $q_i^{k,l-1} \leq q < q_i^{k,l}$. See Figure 1. We also assume that the bid prices are at or above cost, which means that $c_i^{l'} \leq p_i^{k,l}$ for all l' where $(q_i^{k,l-1} \leq q_i^{l'} < q_i^{k,l})$. We say a player is *on a vertical* of its offer (cost) curve at the equilibrium in episode k when its quantity at the trial equilibrium satisfies $Q_i^k = q_i^{k,l}$ for some l . (Another way of putting this is that a player is on a vertical if the market buys all of the quantity it offers on step l and none on step $l + 1$.) We say it is *on a horizontal* of player i when $P_i^k = p_i^{k,l}$ for some l . Note that both prices and quantities are changing. The trial equilibrium either takes its price from some step on the combined offer curve or has its quantity equal to the sum $q_1^l + q_2^{l'}$ for some l and l' . For a player at the corner, either price or quantity can be used as the determinant of the trial solution.

The moves that the players make depend on whether the players are on verticals or horizontals. We treat the cases where both players are on verticals for all episodes in an epoch, both on horizontals, and one on a vertical and the other on a horizontal. Due to the built-in randomness there can be a mix of outcomes for the episodes in an epoch and agents can switch from moving towards one equilibrium to moving towards a different one. However, our interest is in showing the directions of movement of the agents and that these

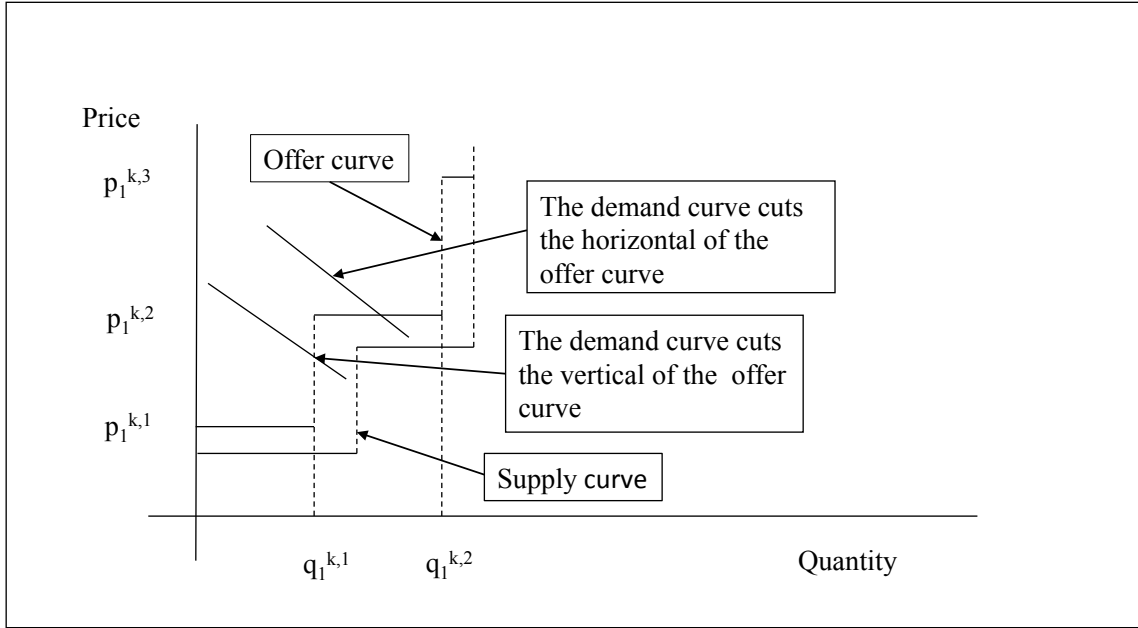


Figure 1: Player 1 definitions.

movements lead the agents toward equilibria that meet known properties. We use the convergence results on the agents to characterize potential equilibria.

5.2 Iterations with Agents on the Verticals of the Offer Curve (Quantities)

In this section we show that the standard optimality conditions for a game determine the moves of the agents and they find an equilibrium that satisfies the Cournot conditions. Since the choices of ϵ and δ are arbitrary, our proofs of agent behavior are also a theoretical characterization of local equilibria that satisfy the Cournot conditions, where we define a local Cournot, Bertrand, or mixed-condition equilibrium to be a point that satisfies the conditions for the Cournot, Bertrand, or mixed-condition equilibrium respectively and neither player can do better than this point in a neighborhood around the point. Thus, we use our proofs about the agents to characterize the different pure strategy equilibria of the oligopoly game with step function supply curves. We begin by showing that the effect of the randomness of the sign of δ does not send the agents off in the wrong direction when there are enough episodes in an epoch. That is, if the marginal revenue is sufficiently positive, player i increases production and vice versa. Note that in the proofs we use a fixed δ with the sign randomly chosen, not a δ drawn from a uniform distribution, as is done in the PROBE AND ADJUST procedure, Figure 19. This greatly simplifies the algebra in the proofs. In our experiments using a fixed δ versus the uniform distribution has no impact on the final solution.

Lemma 1 *Assume both players have trial equilibria on the verticals of their offer curves for every episode*

in an epoch. Let c_i^l be the cost of the step of the cost curve that Q_i^k intersects. As the epoch length $J \rightarrow \infty$, with probability 1 at epoch k , if

$$a - b(2Q_i^k + Q_{-i}^k) - c_i^l > b\delta > 0, \quad (1)$$

then $Q_i^{k+1} = Q_i^k + \epsilon$. If

$$a - b(2Q_i^k + Q_{-i}^k) - c_i^l < -b\delta, \quad (2)$$

then $Q_i^{k+1} = Q_i^k - \epsilon$.

(Proof of the lemma is in §B.1 of the supplementary materials.)

We now show that once the players arrive in the neighborhood of a Cournot solution, they stay in that neighborhood. This does not necessarily happen with the agents in the simulations because a horizontal of an offer curve can randomly set the price with a finite epoch length. We assume that this does not happen in this section.

Lemma 2 *Assume both players have trial equilibria on the verticals of their offer curves for every episode in an epoch. For player 1 at k , let*

$$|a - b(2Q_1^k + Q_2^k) - c_1^l| < 3b\epsilon \quad (3)$$

where c_1^l is the cost coefficient on cost-curve step l . Then

$$|a - b(2Q_1^{k'} + Q_2^{k'}) - c_1^l| < 3b\epsilon \quad (4)$$

with probability 1 as $J \rightarrow \infty$ for all $k' > k$.

(Proof of the lemma is in §B.2 of the supplementary materials.)

From the two lemmas we see that if (1) holds, the Q_i^k changes monotonically until (3) is satisfied. Once (3) is satisfied, it is satisfied with probability 1 in all subsequent iterations. Let (Q_1^*, Q_2^*) solve $a - b(2Q_i + Q_{-i}) - c_i^l = 0$, $i = 1, 2$. This solution can be considered a Cournot solution. We conclude in the following theorem that the players move to the Cournot equilibrium.

Theorem 1 *Assume both players have trial equilibria on the verticals of their offer curves for every episode in every epoch. As $J \rightarrow \infty$, with probability 1 there is a k with a trial solution in the square defined by $(Q_1^* \pm \epsilon, Q_2^* \pm \epsilon)$ and with probability 1, for every $k' > k$ the trial solution at iteration k' falls within the same square.*

(Proof of the theorem is in §B.3 of the supplementary materials.)

Corollary 1 *A solution characterized by the Cournot equilibrium conditions is a local supply curve equilibrium.*

(Proof of the corollary is in §B.4 of the supplementary materials.)

5.3 Iterations when One Player Sets the Price to Block a Supply Step of Another Player

We now examine the case where the total supply of each player is large enough to meet demand when added to the other player's first step. That is, $q_2^1 + q_1^2$ and $q_1^1 + q_2^2$ are large enough so that $a - b(q_2^1 + q_1^2) < c_1^2$ and $a - b(q_1^1 + q_2^2) < c_2^2$. Let P_i^k be the baseline offer price of player i at epoch k that is either setting the equilibrium price or is closest to setting the price. This price is associated with some step l . The offer prices around this price at each episode in epoch k are $P_i^k \pm \delta$. We first show that when the offer prices of both players come close together at some iteration, they then stay close together and the players drive the price down to a classic Bertrand equilibrium. We then show that the offer prices always come close together at some iteration. We present our results using supply curves with two steps to simplify the notation. Since we are looking at local convergence, the proofs remain the same for two adjacent steps where the second step has the same properties of step 2 in the proofs.

Lemma 3 *Assume both players are on the horizontals of their supply curves and $q_2^1 + q_1^2$ and $q_1^1 + q_2^2$ are large enough to satisfy demand at the cost for step 2 and the price on the demand curve at the sums of the first steps of both players is above the costs of the first steps. Assume $|P_1^k - P_2^k| < 2\delta$ and $P_i^k > c_i^1 + \delta$ for $i = 1, 2$. If $|P_1^k - P_2^k| < 2\delta$, then with probability 1 as the epoch length increases to infinity, $P_i^{k+1} = P_i^k - \epsilon$ and $|P_1^{k+1} - P_2^{k+1}| < 2\delta$.*

(Proof of the lemma is in §B.5 of the supplementary materials.)

We now show that $|P_1^k - P_2^k| < 2\delta$ is met starting at some iteration k . Let P_i^M be the profit-maximizing price that player i would charge with $Q_{-i}^k = q_{-i}^1$. Since we are focusing on the interesting cases, assume $P_i^M > \max(c_1^2, c_2^2)$.

Lemma 4 *Assume both players are on the horizontals of their supply curves. Without loss of generality assume*

$$P_1^k < P_2^k - 2\delta. \quad (5)$$

As the epoch length increases to infinity, with probability 1 at some iteration k , $|P_1^k - P_2^k| < 2\delta$.

(Proof of the lemma is in §B.6 of the supplementary materials.)

Theorem 2 *Assume both players have enough capacity, q_i^2 , to meet demand with the addition of their second steps, and for all k , prices in the trial solutions are set by a horizontal step of one of the players. Let $c_1^2 \leq c_2^2$*

and assume that $P_i^M > \max(c_1^2, c_2^2)$. Then the agents move to the Bertrand equilibrium of $c_2^2 - b\epsilon$, that is, there is a k such that for all $k' \geq k$, $c_2^2 > P_1^{k'} + \delta$ with probability 1 as the epoch length increases to infinity.

(Proof of the theorem is in §B.7 of the supplementary materials.)

Corollary 2 *If it exists, a Bertrand equilibrium is a local equilibrium.*

Proof: The result follows directly from Theorem 2. \square

5.4 Iterations with One Player setting the Trial Market Price by its Offer Price and the Other on the Vertical of Its Offer Curve

We now examine the case with one player on the horizontal segment of a step for all iterations while the other is on a vertical segment. We show that with one set of assumptions that is reasonable, one of the agents reaches the quantity breakpoint of one of the steps on its marginal cost curve in a finite number of iterations. For our analysis we assume that the player on the horizontal changes prices by the amount $b\epsilon$. This makes demand change by ϵ . Our first result is to show that the equilibrium condition of the player on the horizontal in this game is the Cournot condition for an equilibrium. However, as long as the price stays above the costs of the relevant steps of the player on the vertical of its offer curve, with the price set by the other player, the player on the vertical sees an increase in profit whenever it increases production because the player setting the price reduces production by the same amount as the the player on the vertical increases production.

Consequently, the price seen by this player does not change and the optimality condition for the player on the vertical is to set price equal to marginal cost. This drives the solution to the vertical of one of the player's marginal cost curves. That is, the player on the horizontal faces decreasing quantities until it hits the lower bound of its supply step or the other player increases the vertical of its offer curve to the vertical on its marginal cost curve. This is an anomalous result in that the player that sees the effect of adjusting quantities, as in Cournot games, actually has the Bertrand equilibrium condition and vice versa.

Without loss of generality, let player 1 be on the horizontal portion of its offer curve. For expository purposes, let this be over the second step of the cost curve for player 1 with an underlying cost of c_1^2 . Let player 2 be on a vertical. To simplify our discussion, let the vertical of its offer curve cut the first step of its cost curve.

Lemma 5 *Assume in all episodes player 1 is on a the horizontal of its offer curve and sets the market-clearing price and player 2 has a quantity on the vertical of its offer curve, the optimality condition for player 1 is the Cournot optimality condition.*

(Proof of the lemma is in §B.8 of the supplementary materials.)

We now show the directions in which the players move. We do this for a two-step supply curve, noting that by changing the indexing this result applies in the general case.

Lemma 6 *Assume player 1 sets the trial price and player 2 is on a vertical of its offer curve for all episodes in epoch k . If $P_1^k > c_2^1$ and $Q_2^k < q_2^1$, where c_2^1 is the cost of the first step of the cost curve for player 2, $Q_2^{k+1} = Q_2^k + \epsilon$ with probability 1 as $J \rightarrow \infty$.*

(Proof of the lemma is in §B.9 of the supplementary materials.)

We now show that once player 1 is in the neighborhood where the Cournot condition holds, it stays there at subsequent iterations with probability 1 as long as player 1 is on a horizontal and player 2 is on a vertical in the trial solutions.

Lemma 7 *Assume that for every episode in every epoch player 1 is on a horizontal of its offer curve and player 2 is on a vertical in all epochs k such that $k' \leq k \leq k''$. At the end of epoch k' , if*

$$|a - 2bQ_1^{k'} - bQ_2^{k'} - c_1^2| < 3b\epsilon, \quad (6)$$

then for all k , $k' \leq k \leq k''$ (6) holds with probability 1 as the epoch length increases to infinity.

(Proof of the lemma is in §B.10 of the supplementary materials.)

Note that this Lemma holds whether or not $Q_2^k = q_2^1$. We now show that when player 1 is far from the Cournot equilibrium condition, it moves monotonically towards satisfying that condition.

Lemma 8 *Assume player 1 is on a horizontal of its offer curve and player 2 is on a vertical in the trial solutions for all episodes in epoch k , $k' \leq k \leq k''$. If*

$$a - 2bQ_1^{k''} - bQ_2^{k''} - c_1^2 > 3b\epsilon, \quad (7)$$

or

$$a - 2bQ_1^{k'} - bQ_2^{k'} - c_1^2 < -3b\epsilon, \quad (8)$$

then, as the epoch length increases to infinity, with probability 1 the left hand side of (7) is monotonically decreasing until the inequality is reversed and the left hand side of (8) is monotonically increasing until the inequality is reversed.

(Proof of the lemma is in §B.11 of the supplementary materials.)

We now show that under reasonable assumptions one player is forced to a vertical of its marginal-cost curve.

Lemma 9 *Assume player 1 is on the horizontal of its offer curve and player 2 is on a vertical for all episodes in epoch k and a sequence of epochs starting from epoch k . Also assume that $c_1^2 > c_2^1 > c_1^1$. Then with probability 1 as the epoch length increases to infinity, this sequence terminates at some epoch k' with one of two conditions holding*

- (i) $Q_2^{k'} \geq q_2^1 - \epsilon$
- (ii) $Q_1^{k'} \leq q_1^1 + \epsilon$

(Proof of the lemma is in §B.12 of the supplementary materials.)

The consequence of this lemma is that the iterations lead to a change in the supply steps that are active. What we see the agents doing in the simulations where this happens is that one player is driven to the edge of a supply step as in the above lemma. Say player 2 reaches the capacity of step 1. Then player 2 sets the price above c_2^1 at a level that satisfies the Cournot equilibrium condition. Figure 2, illustrates this outcome.




Figure 2: An equilibrium where the first steps of both players are utilized and player 1 raises its offer price to satisfy its Cournot equilibrium condition and the price P_1^2 is below p_2^2 .

Corollary 3 *A solution characterized by player i satisfying the Cournot equilibrium condition and player $-i$ at a quantity q_{-i}^l can be a local equilibrium.*

Proof: The result follows from Lemmas 8 and 9. \square

We refer to this equilibrium as a mixed-condition equilibrium because one player satisfies the Cournot condition and the other the Bertrand condition.

6 Existence and Uniqueness

In this section we show that an equilibrium always exists and there can be multiple equilibria. Since we can characterize the nature of the equilibria, it is relatively easy to find them in the case of a duopoly. Let

the Cournot local equilibrium be denoted by (P^c, Q_1^c, Q_2^c) and the Bertrand and mixed-condition equilibria be denoted by superscripts b and m respectively. We make a behavioral assumption in the spirit of the Cournot and Bertrand conjectures to determine what the players optimize: assume that each player takes the offer curve of the other player as fixed in optimizing its bid. We begin by showing that at least one local equilibrium exists, the Cournot solution.

Theorem 3 *A local Cournot equilibrium exists and it is the only Cournot equilibrium.*

(Proof of the Theorem is in §B.13 of the supplementary materials.)

To understand why the solution is not necessarily unique, we first examine the relationship between the Bertrand and mixed-condition equilibria. In our analysis of the Bertrand equilibrium, we assume that the player with the lower-cost second step sets the total quantity and price. Let the player with the lower cost be i . Assume that $c_{-i}^1 < c_i^2$. Then in the Bertrand equilibrium $Q_{-i}^b = q_{-i}^1$. Our results for the mixed equilibrium are developed for $Q_{-i}^m = q_{-i}^1$ as well.

Lemma 10 *At the Bertrand equilibrium, if it exists,*

$$a - 2bQ_i^b - bQ_{-i}^b - c_i^2 \leq 0. \quad (9)$$

or, equivalently,

$$c_{-i}^2 - bQ_i^b - c_i^2 \leq 0 \quad (10)$$

Proof: Since $Q_{-i}^b = q_{-i}^1$ for $P^b = c_{-i}^2$, player $-i$'s quantity is fixed and player i sees the Cournot equilibrium condition on its profits. If

$$a - 2bQ_i^b - Q_{-i}^b - c_i^2 > 0, \quad (11)$$

then player i can increase its profits by lowering P^b and increasing Q_i^b and the Bertrand solution is not a local equilibrium. The condition (10) follows from noting that the limiting price is $c_{-i}^2 = a - bQ_i^b - bQ_{-i}^b$. Thus, the result holds. \square

If (11) holds, then we have a mixed-condition local equilibrium with $P^m < c_{-i}^2$. When the Cournot condition equals 0 in the Bertrand equilibrium, the mixed-condition and Bertrand equilibria coincide. If (9) holds as a strict inequality, it is possible to have a mixed-condition equilibrium with $Q_{-i}^m = q_{-i}^2$ because it might be more profitable for player i to reduce production and let the second step for player $-i$ be positive.

Whichever equilibrium obtains, the player setting the price, player i , has its actions determined by the Cournot condition in both the Bertrand and mixed-condition local equilibria. Furthermore, in both the Bertrand and mixed-condition equilibria, the other player, player $-i$, sees that marginal revenue equals price and $P^b \geq c_{-i}^l$ or $P^m \geq c_{-i}^l$ when the quantity in the solution is q_{-i}^l . (We use l here because it is possible for other steps than the first to be the Bertrand quantity as in our examples.)

In the Bertrand game, if player $-i$'s offer curve is its marginal cost curve, which is the optimal offer curve when marginal revenue equals price, then any attempt by player i to increase the price either by raising the price on its offer curve above c_{-i}^2 or lowering the quantity on its vertical below Q_i^b by a small amount leads to the price remaining at c_{-i}^2 because of increased production by $-i$. Thus, there is no move by player i that can increase profits and there is the potential for multiple equilibria in this case. Note that the Bertrand equilibrium might not exist when there is insufficient capacity for the price-setting player to meet demand. Here the Bertrand equilibrium is replaced by Edgeworth cycles, which we observe in the simulations.

When the profits for player i are lower in the Bertrand game than the Cournot game, player i can be trapped in this equilibrium. However, when the profits are lower in the Cournot game, player i can move from the Cournot solution to the Bertrand solution simply by setting its price to c_{-i}^2 and its quantity to its maximum capacity. The reason this is a profitable move is that the offer curve of player $-i$ has to be at or above its marginal cost curve and the profit for i with $-i$ using an offer curve of above its marginal cost curve is at least as large as with $-i$'s marginal cost curve.

With the mixed-condition equilibrium, the profit of player $-i$ decreases for any restriction in quantity below Q_i^m , even with the price below c_{-i}^2 , because the Cournot condition is satisfied at that equilibrium for player i . As in the Bertrand game, when player $-i$'s offer curve is its marginal cost curve, trying to raise prices above c_{-i}^2 does not improve profits for player i . The Cournot equilibrium has a higher profit than the mixed-condition equilibrium for player i simply because the Bertrand condition for player $-i$ leads to a quantity above the Cournot quantity for player $-i$, given the price set by player i . However, given the quantity of player $-i$ in the mixed-condition equilibrium, player i decreases its profits when moving to the Cournot quantity. After the next section, we provide an example with multiple equilibria.

We have characterized the equilibria when both players are on horizontals or verticals of their offer curves and one player is on a horizontal and the other is on a vertical. These cases encompass all possible pure-strategy equilibria. In the case of the mixed-condition equilibria, there are two sets of equilibria that have to be considered, since it is possible for both players to set prices. Thus, one can examine equilibria of these forms to characterize all of the local equilibria in the game and test to see if these local pure-strategy equilibria are true equilibria. This means that it is not insurmountable to enumerate the multiple pure-strategy equilibria when they exist.

Our results show far fewer equilibria than is the case with continuous supply-function equilibria. The reason for this is simple. Step functions have two slopes, 0 or infinite and the best responses to those slopes are the Cournot or Bertrand player optimizations. Continuous offer curves have an infinite number slopes, leading to an infinite number of best responses and equilibria.

7 Exploring the Analytic Results with the Computational Model

Our computational model, which we sketched in §4, is embodied in the NetLogo program `SupplyCurveBidding.nlogo`, which may be accessed at <http://asadai.wharton.upenn.edu/nlogo/SupplyCurveBidding.html> and is described in §C of the supplementary materials. In this and the next two sections section we report on experiments done with our computational model. §D contains further supplemental information derived from `SupplyCurveBidding.nlogo`.

The analytic results, above, identify equilibria for three distinct cases:

1. All players are on their verticals (Theorem 1), in which case we have a Cournot equilibrium.
2. The players have supply much in excess of demand (given costs), in which case we have a Bertrand equilibrium (Theorem 2) and the players compete away their profits.
3. The mixed-condition equilibrium case, with one player on the horizontal (the demand curve intersects the player’s bid curve) and the other player(s) on the vertical (§5.4).

We note that the equilibria so characterized are one-shot equilibria in what is actually an iterated game. Our analytic results demonstrate that in the limit agents playing iteratively with `PROBE AND ADJUST` and using the `Own Returns` criterion can reach these one-shot equilibria and settle there.

In this section we present and explore a number of exemplary scenarios in which two firms, each with two plants, learn to bid supply curves in the market we have described. The upshot of this discussion is that finite computational agents using `PROBE AND ADJUST` and acting myopically do robustly learn to reach the equilibria analytically characterized. In the following section, we also briefly address sensitivities of the model, which in our experience is quite robust.

In the interests of brevity and of communicating the main results, our presentation here is compressed. The appendix sections contain supplementary information about the behavior of our simulation model and we have made our model available online for anyone to explore.




Figure 3: “Supply & Demand Curves” chart with the Base Case Parameterization. From Figure 20 in the supplemental materials

7.1 On the Verticals

Figure 3 displays the “Supply & Demand Curves” chart at the conclusion of a typical run with the Base Case Parameterization. Notice that the agents have sorted themselves out so as to play on the verticals: the demand curve intersects the supply curve at a vertical. The initial configuration of each agent’s $(2N - 1) = 5$ bid parameters puts it close to the final configuration shown here. For example firm 1’s initial base prices for plants 1, 2, and 3 are 16, 150, and 150, while at the conclusion of the run we see that these base prices have evolved to 15.4, 145.8, and 152.6 (see widget monitors in the lower left corner region of the interface). In this Base Case example, the agents do not discover an equilibrium so much as go to it, having been placed nearby.

The conditions of Theorem 1 apply in the Base Case. With $a = 200$, $b = 1$, and $c_1 = c_2 = 15$, solving $a - b(2Q_i + Q_{-i}) - c_i = 0$ yields $Q_1^* = Q_2^* = 61.67$, so the total quantity offered at equilibrium would be 123.33, which compares well with 123.2 obtained as an average over 30 replications of the Base Case Parameterization. (See Table 3 in the supplemental materials). Briefly, these runs tell us that when the firms both begin with prices well above this equilibrium, our agents are able to find it while continuing to explore. Modest changes in the firms’ initial base prices for plants 1, 2, and 3 do not affect the outcome. Figure 4 shows the Supply & Demand Curves plot from representative run with somewhat lower initial prices. The same results obtain as in Figure 20. Note, however, that the demand curve runs between bid quantities of the two firms; firm 1 to the left, firm 2 to the right.




Figure 4: P2BothIs25 Parameterization, supply and demand plot for a representative run (both firms update with Own Returns), with both firms starting with high bids. From Figure 29 in §D.

7.2 Bertrand Competition

The Bertrand250Both parameterization (available as a predefined case in the program) differs from the Base Case in the quantities available. $\text{MaxQ2Firm1} = \text{MaxQ2Firm2} = 250$ (instead of 100). Thus the total quantity available is 540 and quite ample for the demand, thereby meeting the condition for our Theorem 2 results.

Table 4 in the supplementary material summarizes the results, which are pictured vividly in Figure 5. The firms’s costs for their second plants are 15 and the market price is just slightly above that.



Figure 5: Bertrand250Both Parameterization, prices for a representative run (both suppliers update with Own Returns). From Figure 21 in §D.

The initially high prices result from starting quantities on the steps of the low-price player, which we term player 1, that do not meet all demand. The market takes all of either the first or first two steps of this player. At the same time the player 2 finds it profitable to keep its price above the price set by the player 1, searching for a mixed equilibrium. Since it is profitable for player 1 to increase the size of the steps taken by the market, eventually it is more profitable for player 2 to lower prices and compete as in the Bertrand game. The second bump in prices reflects a stochastic version of an Edgeworth cycle. Through a sequence of random outcomes, player 2 shortens its lower-priced steps below what is needed for this player to meet demand, while having higher prices than the other player on these steps. Once the prices of these steps randomly drop below the prices on the steps of player 1 that are taken by the market, player 1 finds it profitable to raise prices, aiming for the mixed solution. Player 2 eventually expands the winning supply steps to the point where player 1 reverts to playing Bertrand and profits evaporate.

7.3 Mixed-Condition Equilibrium

We use the P2BothIs25 built-in case to illustrate the mixed-condition equilibrium described in §5.4. The P2BothIs25 parameterization differs from the Base Case only in the initial base prices for bid quantity 2 for the firms. $\text{InitialBasePrice2Firm1} = \text{InitialBasePrice2Firm2} = 25$. In this setting one firm reliably settles on its verticals and the other on its horizontal, yielding very different returns for the players. The firms take their positions randomly, since they start out identically. Because of this, summary statistics on multiple runs hide information. Thus, we hold the discussion in terms of a single, quite representative run.

The initial profit decline in the figure comes from the firms engaging in Bertrand price cutting until one player finds itself on a vertical and the other on a horizontal. Then they move to the mixed-condition equilibrium.

Figures 6 and 7 show results from a common run. We see from Figure 7 that firm 1 is on the horizontal




Figure 6: PSBothIs25 Parameterization, profits for a representative run (both firms update with Own Returns). See Figure 24 in §D for another example.




Figure 7: PSBothIs25 Parameterization, supply and demand plot for a representative run (both firms update with Own Returns). See Figure 24 in §D for another example.

and firm 2 is on the vertical. The mean running total quantity is 153.96. Firm 2 has its entire capacity of 120 in play, leaving firm 2 with a capacity of 33.96. From expression (24) in §5.4 we expect firm 1's quantity to be

$$\frac{(200 - 120 - 15)}{2} = 32.5$$

The computational agents have indeed arrived close to the predicted outcome. The firm on the vertical is earning profits of about 3936 at the end of the run(s) and firm on the horizontal about 1283.

8 Further Explorations with the Computational Model

There is a useful comparison to be made between simulated annealing, as standardly deployed for parametric (non-strategic) optimization, and PROBE AND ADJUST deployed as we have described it for finding equilibria of games. Both simulated annealing (Johnson et al., 1989) and PROBE AND ADJUST are randomized, meliorating search meta-heuristics. They are both hill climbers with a degree of tolerance for local ups and downs. A given optimization problem may have multiple optima; different paths leading to an optimum may present differing challenges for a heuristic searcher. For these reasons it is appropriate to use simulated annealing by conducting many replications with diversified initialization. It is well known that simulated annealing is often very successful under these conditions and, therefore, has been widely adopted.

As in optimization, strategic interaction problems may have multiple equilibria and the paths to an equilibrium may be fraught for a heuristic player. For these reasons it is appropriate to use PROBE AND ADJUST by conducting many replications with diversified initialization. Our purpose in this section is to report on a number of such experiments, tied to our present context.

To begin, we observe that the Base Case Parameterization, discussed in §7.1, and the P2BothIs25 built-in case, discussed in §7.3, are identical except for the initial base prices for the firms' second step. In the Base Case, the firms start high and PROBE AND ADJUST leads them to an equilibrium on the verticals. In P2BothIs25 both firms start low and are led to a mixed equilibrium. What happens if one firm starts high and the other starts low? Reliably, we get the equilibrium with both firms on their verticals.

8.1 P2Both2Is25 with 6 Bid Steps

We saw in §7.3 that the P2Both2Is25 setup leads reliably to a mixed equilibrium condition, with one firm on its verticals and the other on its horizontal.

The built-in scenario P2Both2Is25 6 Steps is a setup like P2Both2Is25 but with each firm having six, rather than three, bid steps. Figures 8 and 9 are from a representative run in which both firms begin by bidding low, 16 for each bid step. We see from Figure 8 that the resulting market is quite volatile, although by the end of the run markets have at least temporarily stabilized somewhat. The reason for the volatility is that the prices and quantities that do not affect profits move in a random walk. This results in random occurrences of Bertrand behaviors from which the firms recover and move towards the mixed equilibrium.

Firm 1 has at this point an average running profit of 4427, firm 2, 2202. Compare this to the results in §7.3. In the present case, firm 2 is on the horizontal, firm 1 on the vertical (Figure 10). Compared to the §7.3 results, however, the disadvantaged firm, on its horizontal, is doing much better and the industry as a whole is more profitable.



Figure 8: P2BothIs25 parameterization, but with 6 bid steps for each firm, prices for a representative run (both suppliers update with Own Returns), each firm initializing with low bids. From Figure 28 in §D.

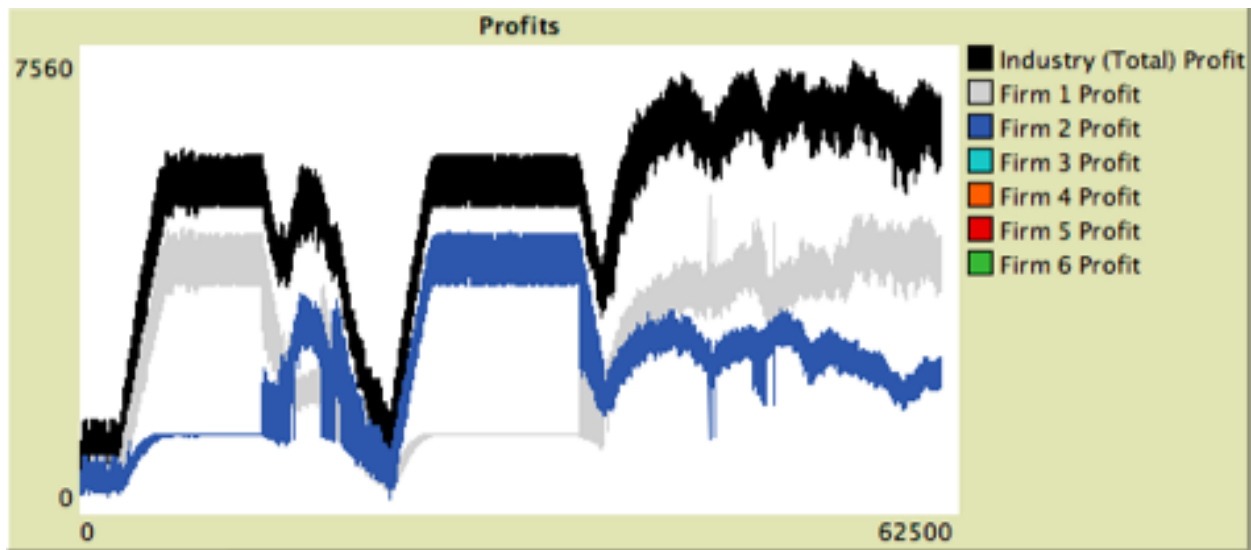


Figure 9: P2BothIs25 parameterization, but with 6 bid steps for each firm, profits for a representative run (both suppliers update with Own Returns), each firm initializing with low bids. From Figure 28 in §D.

If both firms, having 6 steps each, initialize with 100 as their prices, we see a Bertrand-style race to the bottom, followed by a volatile recovery, punctuated by Bertrand-like price crashes, but roughly settling as in the previous case. We see that the ability to ultimately find the equilibrium is unaffected by the number of steps in the offer curves. However, our agent is not smart enough to override the random changes that disrupt the equilibrium. See Figure 30 in §D for a representative run.

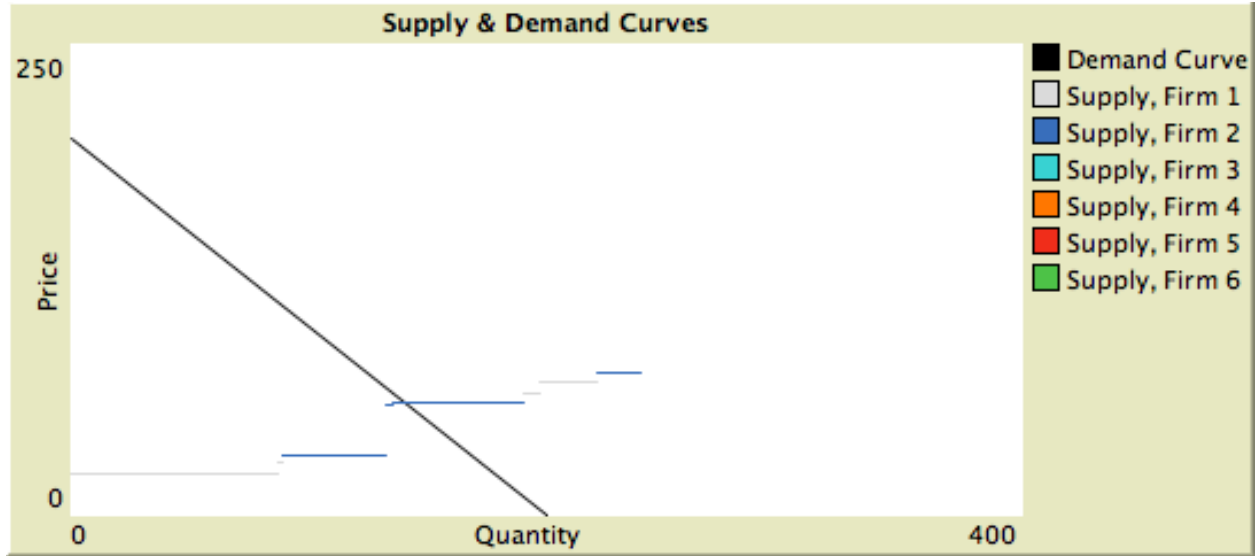


Figure 10: P2BothIs25 parameterization, but with 6 bid steps for each firm, supply and demand curves for a representative run (both suppliers update with Own Returns). From Figure 28 in §D.

8.2 Q2Firm1Is250

In the Bertrand250Both parameterization, discussed in §7.2, both firms have an excess of capacity given the demand function, leading them to behave in classic Bertrand fashion and to compete away their profits. The Q2Firm1Is250 built-in parameterization differs from the Bertrand250Both setup only in that firm 2's capacity for its second plant is 100, rather than 250. In consequence, firm 2 reliably is on the vertical and firm 1 on the horizontal, yielding a persistent advantage to firm 2. (Table 5 in the supplementary materials summarizes results from 30 replications.)

With the Q2Firm1Is250 setup, but with six bid steps for each firm and prices starting low (at 16 for each step), results are variable and the market is at times unstable. In a typical run firm 2 is usually on its verticals and has a substantial advantage (Figure 31 in §D shows results from a representative run). In some runs, however, firm 1 manages to switch positions with firm 2. If firm 2 foolishly starts out with high prices, firm 1 is effectively in the position of a monopolist. As firm 1 learns this, however, firm 2's prices drift lower until they encounter firm 1's rising prices. A chaotic sorting out ensues with firm 2 reliably achieving its position on the verticals.

8.3 Bertrand250Both 6 Steps

The Bertrand250Both 6 Steps built-in setup is that of Bertrand250Both, discussed in §7.2, but with six bid steps per firm instead of three. Comparing the typical results shown in Figure 12 with the results for the Bertrand250Both case we find that the move to six bid steps creates a smoother-functioning market, from the firms' perspective, and leads to a small degree of rent (the running average profits are about double to

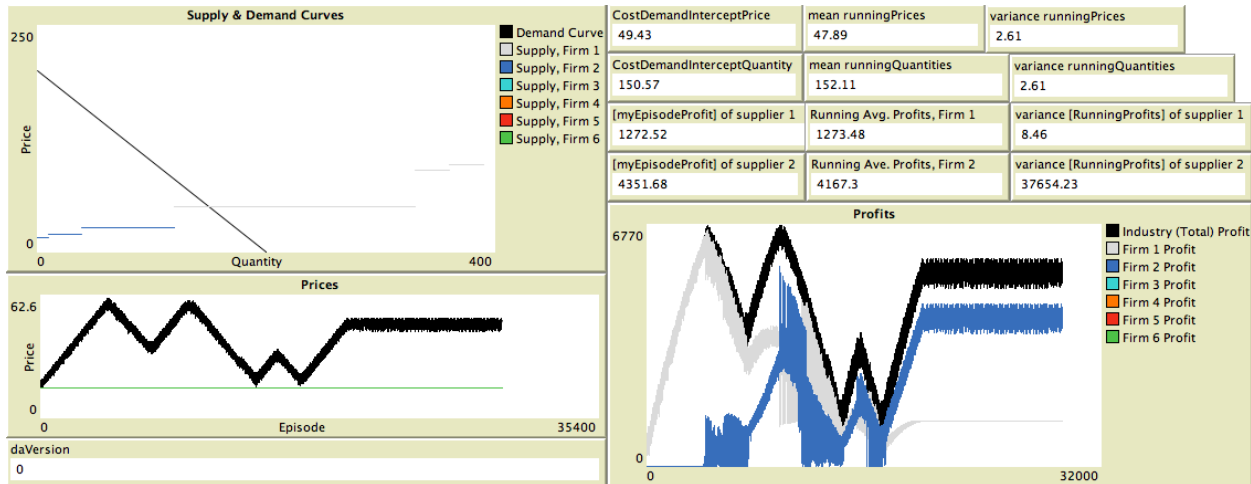


Figure 11: P2BothIs25 parameterization, but with 6 bid steps for each firm, supply and demand curves for a representative run (both suppliers update with Own Returns), firm1 initialized to 16 for prices, firm 2 to 100. From Figure 33 in §D.

those of the Bertrand250Both case).

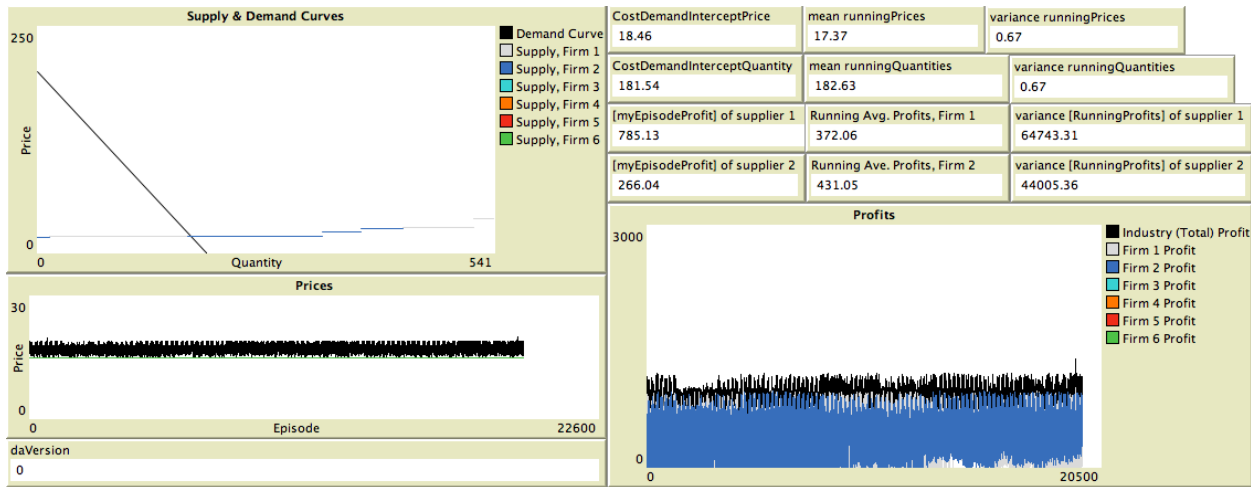


Figure 12: Bertrand250Both 6 Steps setup, representative run. Each firm has 6 bid steps. All firms start by bidding low. From Figure 35 in §D.

The results consist of oscillations in the prices and profits. The added steps lead to repeated Edgeworth behaviors because there are more steps having the random adjustments that regularly trigger the price increases that occur with the bumps in the Bertrand simulations having three steps.

8.4 Multiple Equilibria

In the theoretical results we show the potential for multiple equilibria. Multiple equilibria appear in the simulations as well. With the base-case parameters and a capacity of 250 on the second supply step, the players find the Bertrand solution because the players start with a very wide first step, which means the offer curves initially cut the demand curve on the horizontals and the players compete by lowering prices. When the initial offer curves of both players are combined and cross the demand curve at a vertical, then the players find the Cournot equilibrium. We make this happen by reducing the size of the first step in the initial offer curve of each player to 34 and price the second step at 90. If player 1 has a high initial price of 90 on the second step while player 2 has a much lower price of 35, the resulting solution is the mixed-condition equilibrium. If the second step of player 2 is wide enough, the mixed equilibrium is a true equilibrium in that neither player can do better. Thus, using different starting points, the players find three different equilibria, verifying that competing using supply curves leads to multiple equilibria.

9 Exploring a Non-Myopic Strategy: MR-COR

In all of the runs discussed so far all of the firms are updating with the Own Returns policy. Under Own Returns after each episode of play, the firms individually record their returns categorized by bid segment and whether the bid is above or below the base quantity. Specifically, if there are N bid segments there are $(N - 1)$ bid breakpoints. For the resulting $(2N - 1)$ values for prices and quantities, the agent records its returns in three bins: one for when the bid is very near the baseline, one for returns when the bid amount is above the baseline, and one for when the bid amount is below. Then, at the end of its epoch (typically 100 episodes), the agent calculates the most profitable directions to change the prices and quantities and adjusts them in that direction. Under the Market Returns update policy, the procedure is the same, except that the agent records and looks to the returns in its industry, which we label the market returns, and so is an altruist. Under the MR-COR policy (Market Returns, Constrained by Own Returns), the agent keeps a double set of bins, one for its own returns and one for the market returns. When the agent goes to update its prices and quantities it compares its performance during the epoch with that of the rest of the firms in the market. If the agent's returns are only slightly below or better than its target share of market's returns, then it updates with Market Returns; otherwise it uses a trigger strategy and updates with Own Returns. Our purpose in this section is to report on MR-COR, comparing it to Own Returns.

9.1 Beyond on the Verticals

The material in this subsection complements that in §7.1 and the Base Case parameterization, but with the firms' update policies changed.

With both players updating with MR-COR, the players are able to achieve tacit collusion at the expense of the customers. (See Table 6 in the supplementary materials.) A single firm acting as a monopolist with the




Figure 13: Bertrand250Both Parameterization, prices for a representative run (both suppliers update with MR-COR). From Figure 22 in §D.

combined plants of both players achieves an industry profit of about 9,000. This is indistinguishable from the performance of our (tacitly) colluding MR-COR firms. When one player updates with Own Returns and the other with MR-COR, the selfish player, playing Own Returns, will win out over the MR-COR player, but it would do even better for itself by switching to MR-COR, provided the other player sticks with its MR-COR strategy. The MR-COR player, however, would do better to switch to Own Returns if its counterpart does not switch to MR-COR. (See Table 7 in the supplementary materials.)

9.2 Beyond Bertrand

This section complements §7.2. Both discuss the Bertrand250Both parameterization.

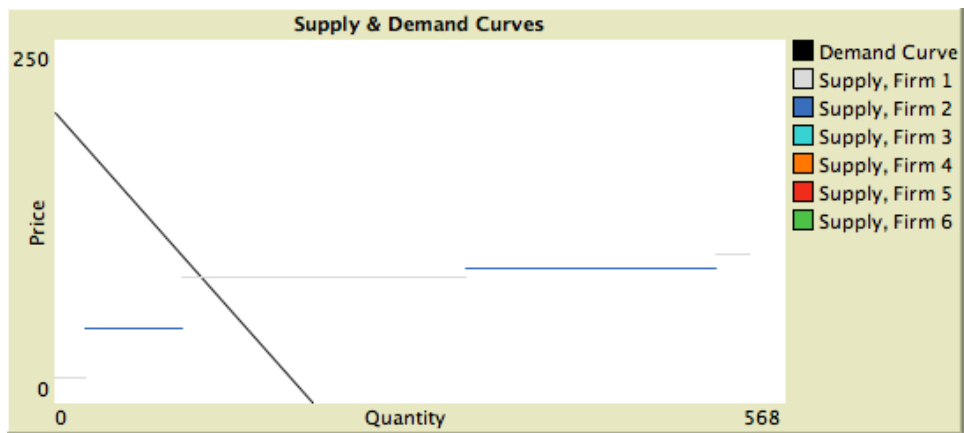


Figure 14: Bertrand250Both Parameterization, supply and demand curve for a representative run (both suppliers update with MR-COR). From Figure 22 in §D.

When both firms update with MR-COR, the results are quite different. See Figures 13 and 14. There

is a noticeable lack of settled stability, but a clear pattern does emerge: the supplier firms jointly achieve profits slightly below the monopoly level, with ongoing jostling. They are both much better off than when both update with Own Returns.

When one firm updates with Own Returns (here firm 1) and the other (firm 2) with (MR-COR), characteristically the Bertrand pattern usually results. Figure 15 displays an alternative pattern we find occasionally. What happens here is that firm 2, updating with MR-COR and receiving returns much less than firm 1, drives the price down towards the Bertrand level, until by chance it gets a return near to or above firm 1's. At this point firm 2 begins updating based on Market Returns, leading to general prosperity for the industry, especially itself with firm 1 lagging. By chance, firm 2 gets on its vertical with firm 1 on its horizontal. Firm 1, looking at its Own Returns, increases its price in the direction of the Cournot quantity for firm 2's current quantity. Firm 2 acts so as to increase the industry's performance and gradually explores shifting quantity to the left, so as to be in the market. Once the basic pattern is set, with firm 2 having substantial production and on its vertical, firm 2 leads a gradual improvement for the industry, with its growth roughly in line with that of firm 1.




Figure 15: Bertrand250Both Parameterization, profits for a representative run (firm 1 updates with Own Returns, firm 2 with MR-COR). From Figure 23 in §D.

9.3 Beyond the Mixed-condition Equilibrium

This section corresponds to §7.3 and discusses the P2BothIs25 built-in case.

When both agents update with MR-COR we expect and we find that they do about equally well. Figures 16 and 17 show representative results. The firms succeed in obtaining monopoly prices and quantities for the industry (just under 9,000) and in splitting the profits more or less equally (at about 4,500 each). Figure 17 is particularly interesting in that it shows both firms on the verticals, with the demand function running between the steps of different firms.

Finally, there is a marked difference if one firm updates with Own Returns and the other with MR-COR.

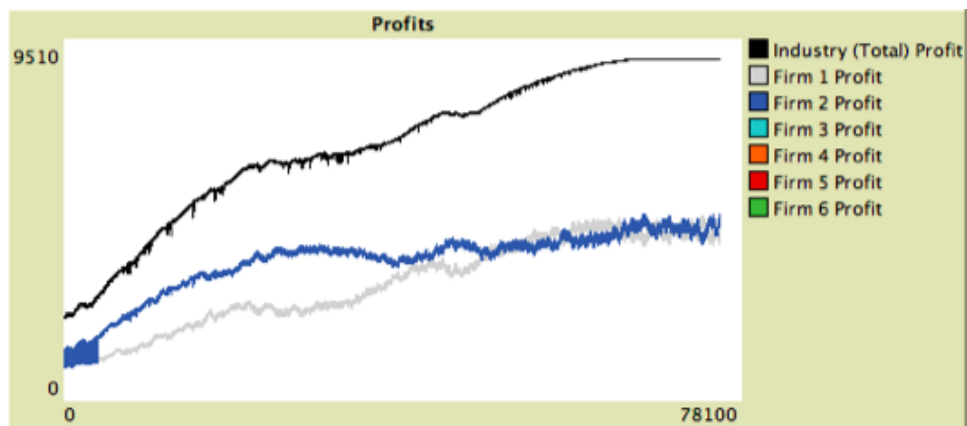


Figure 16: P2BothIs25 Parameterization, profits plot for a representative run (both firms update with MR-COR). From Figure 25 in §D.

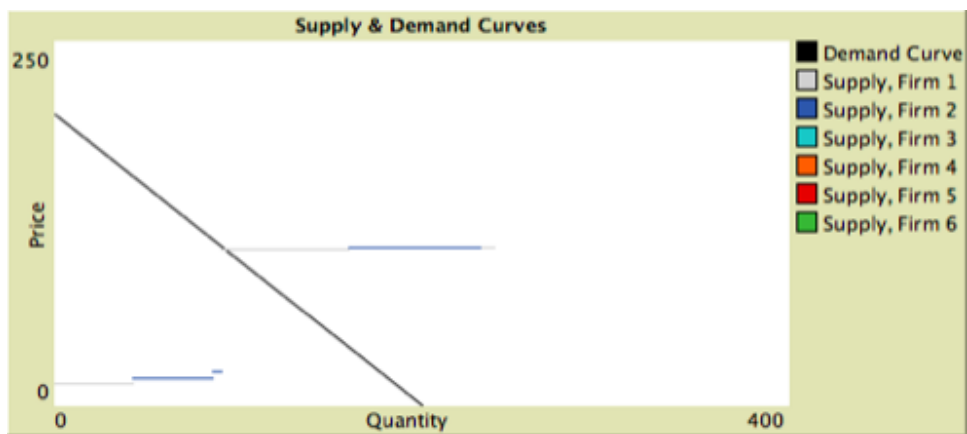


Figure 17: P2BothIs25 Parameterization, supply and demand plot for a representative run (both firms update with MR-COR). From Figure 25 in §D. See also Tables 9 and 10 in the supplemental materials.

The MR-COR player loses out, although the Own Returns player would do better by switching to MR-COR. (See Table 8 in the supplemental materials.)

9.4 Q2Firm1Is250

We saw in §8.2, discussing the Q2Firm1Is250 case, that firm 2 has a reliable advantage over firm 1, whether with three bid steps or six. This advantage remains even with MR-COR as an option.

9.5 NFirmsSimple

In the NFirmsSimple case, the firms are identical. They each have two plants with costs of 0 and capacities of 15 and 55. Their five initial capacity breakpoints are (0.2 0.4 0.6 0.7 0.8) and their six initial prices are all 2. Using five supplier firms and setting QIntercept to 300, when all firms update with MR-COR, they consistently achieve a high level of tacit collusion, driving up prices and profits for all. Figure 18 is from a representative run.

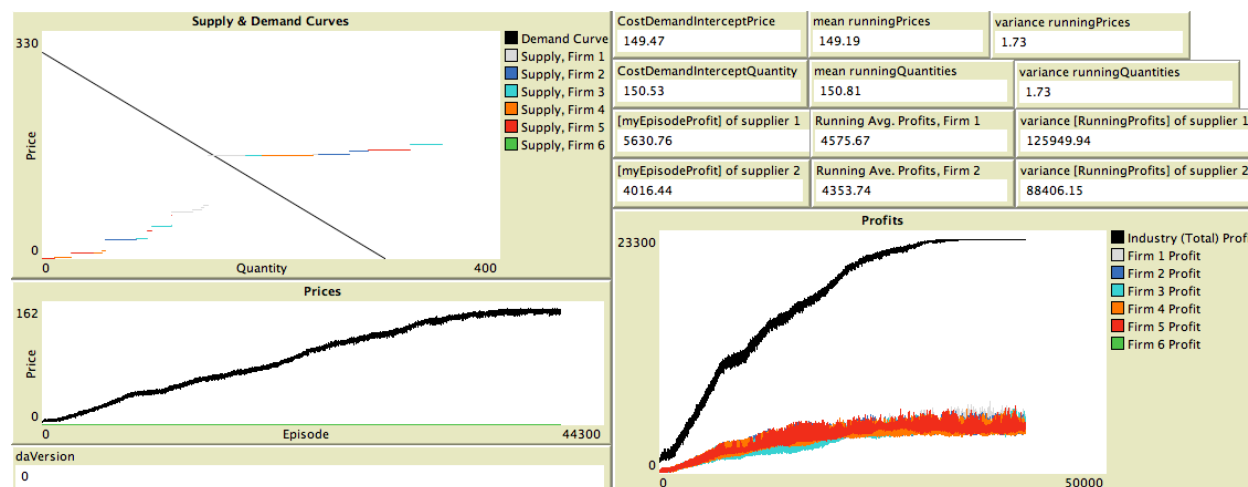


Figure 18: NFirmsSimple parameterization, plots for a representative run. Five firms update with MR-COR; each firm has 6 bid steps and two plants of capacity 15 and 55. All firms start by bidding low. QIntercept is 300. From Figure 34 in §D.

10 Summary and Conclusion

As a summary we offer the following points:

1. We have modeled both analytically and with computational agents a market where firms bid step-function offer curves. An important market of this form is the day-ahead electricity market in which suppliers bid step-function offer curves and these bids are used by an independent system operator to find market prices and quantities.

2. As we have noted, the step-function structure of the offer curves leads to analytic difficulties. Except as noted above, the literature has generally avoided addressing the difficult issues associated with the discontinuities of step-functions.
3. We examine players who use a plausible learning regime in such markets. We call the regime PROBE AND ADJUST.
4. We characterize analytically the behavior of agents using this learning regime in the context of step-function offer curves, setting a template for characterizing the behavior of agents. Without a theoretical grounding for the representations of agent behavior and simulation outcomes in standard economic theory, agent-based modeling risks losing its credibility as a modeling tool if models generate anomalous results through ad hoc simulations.
5. We characterize the potential equilibria using analytical methods and verify that the agents can find these equilibria.
6. The agents exhibit transient phenomena that offer insight into market dynamics.
7. Using the Own Returns update policy in PROBE AND ADJUST, the agents play a repeated (iterated) game and find equilibria of the one-shot (stage) game, as well as the iterated game.
8. Using the MR-COR trigger policy with PROBE AND ADJUST, the agents are often able to achieve and maintain tacit collusion that yields them much higher profits.
9. PROBE AND ADJUST in the context of repeated games involving adjustment of continuous decision variables (price, quantity in our case) may usefully be compared with simulated annealing and other metaheuristics used for standard optimization. In both cases, the computational method is often able to discover otherwise hard to find solutions / outcomes. PROBE AND ADJUST should be seen both as an approximation of a realistic decision procedure by players and as method of finding equilibria and near-equilibria in games. At the same time the agents can be viewed as a plausible simulation of the interactions among players as they move to an equilibrium.

In the conduct of this research we have found a symbiotic relationship between agent-based simulations and analytic results. For example, our agents found the mixed-condition equilibrium which we originally had not considered. We were then able to analytically characterize the behavior of the agents moving to of this equilibrium. Also, some simulations exhibit Edgeworth behaviors that result from incremental learning and adjustment rather than capacity limits. Thus, the agent simulations served as hypothesis discoverers. As a second example, the dynamics of the simulations highlighted the different rôles of price and quantity positions in determining where markets settle. In both the myopic case, with both firms using Own Returns as an update policy and with both firms using MR-COR, the firms are individually and collectively better off if the equilibrium is on the quantity break points of their offer curves (“on the verticals”), rather than the price break points (“on the horizontals”).

MR-COR represents a first cut example of how the model can extend modeling capabilities beyond the standard models. We intend to take this work in several directions. One is to examine agency effects of performance rewards. Here we can look at the effects of rewarding traders and/or CEO's based on total profits versus the profits of each firm versus the other firms. A second is to examine the effects of heterogeneous forecasts of demand on market equilibria. This is useful simply because the most optimistic player overbuilds, undercutting market power. Thirdly, when each player has beliefs about the other player's response to its quantity-setting actions that are consistent with the actual response and costs are linear, the Bertrand equilibrium obtains. It is possible to simulate the process of belief formation and see its effect on the equilibrium.

Agent-based modeling allows us to combine complicating factors in markets in a single model and observe their interactions rather than rely on analytic results in models that do not capture the interactions among the complexities. At the same time the behavior of the agents has to be grounded in theory in the standard cases as part of validating the model, as we have done here so that results are not artifacts of ill-conceived agents.

References

Anderson, E.J. and H. Hu, 2008, Finding supply-function equilibria with asymmetric firms, *Operations Research*, 56(3), 697-711, May-June.

Anderson, E.J. and A.B. Philpott, 2002, Using supply functions for offering generation into an electricity market, *Operations Research*, 50(3), 477-489, May-June.

Anderson, E.J. and H. Xu, 2005a, Supply function equilibrium in electricity spot markets with contracts and price caps, *Journal of Optimization Theory and Applications*, 124(2), 257-283, February.

Anderson, E.J. and H. Xu, 2005b, Optimal bidding in an electricity market with discontinuous market distribution function. Working paper, Australian Graduate School of Management University of New South Wales.

Baldick, R., R. Grant, E Kahn, 2004, Theory and application of linear supply function equilibrium in electricity markets, *Journal of Regulatory Economics*, 25(2), 143-167.

Baldick, R. and W.W. Hogan, 2002, Capacity constrained supply function equilibrium models of electricity markets: Stability, non-decreasing constraints, and function space iterations, working paper, Center for Business and Government Harvard University.

Baldick, R. and W.W. Hogan, 2004, Polynomial approximations and supply function equilibrium stability, working paper JFK School of Government, Harvard University.

Bagnall, A.J. and G.D.Smith, 2005, A multi-agent model of the UK market in electricity generation, *IEEE Transactions on Evolutionary Computation*, 9(5), 522-536, October.

Boone, J., 2002, 'Be nice, unless it pays to fight:' a new theory of price determination with implications

for competitive policy, working paper FS-IV 02-18, Tilburg University.

Brandts, J., P. Pezanis-Christou, and A. Schram, Competition with forward contracts: a laboratory analysis motivated by electricity market design, *Economic Journal* 118(1), 192-214, January.

Bunn, D.W. and F. Oliveira, 2001, Agent-based simulation-An Applications to the New Electricity Trading Arrangements of England and Wales, *IEEE Transactions on Evolutionary Computation*, 5(5), 493-503, October.

Bunn, D.W. and F. Oliveira, 2003, Evaluating individual market power in electricity markets via agent-based simulation, *Annals of Operations Research*, 121, 57-77.

Bunn, D.W. and F.S.Oliveira, 2008, Modeling the Impact of Market Interventions on the Strategic Evolution of Electricity Markets, *Operations Research*, 56(5), 1116-1130, September-October.

Bower, J. and D.W. Bunn, 2000, A model-based comparison of pool and bilateral market mechanisms for electricity trading, *Energy Journal*. 21(3), 1-29.

Contreras, J., O. Candiles, J. Ignacio de la Fuente, T. Gomez, 2002, A cobweb bidding model for competitive electricity markets, *IEEE Transactions on Power Systems*, 17(1), 148-153, February.

Day, C.J., B.F. Hobbs and J.-S. Pang, 2002, Oligopolistic Competition in Power Networks: A Conjectured Supply Function Approach. *IEEE Transactions On Power Systems*, 17(3), 2002, p. 597-606.

Edgeworth, F.Y., 1925, The pure theory of monopoly, in ed. Edgeworth, *Papers Relating to Political Economy*, vol. 1. Burt Franklin, NY.

Erev, I., and A. E. Roth, 1998, Predicting how people play games with unique mixed-strategy equilibria, *American Economic Review* 88, 848-881.

Ernst, D., A. Minoia, M. Ilic, 2004, Market dynamics driven by the decision-making power producers, *Bulk Power System Dynamics and Control - VI*, August 22-27, 2004, Cortina d'Ampezzo, Italy 41-47.

Fabra, N., N. von der Fehr, and D. Harbord, 2006. Designing Electricity Auctions, *RAND Journal of Economics*, 37(1), 23-46.

Green, R.J., 1997, The Electricity Contract Market in England and Wales. *Journal of Industrial Economics*, 47(1), 107-123

Green, R.J. and D.M. Newbery, 1992, Competition in the British electricity spot market, *Journal of Political Economy*, 100(5) 929-953 October.

Hobbs, B.F. and F.A.M. Rijkers, 2004, Strategic generation with conjectured transmission price responses in a mixed transmission pricing system-Part 1: Formulation, *IEEE Transactions on Power Systems*, 19(2), 707-717, May.

Hobbs, B.F., F.A.M. Rijkers, and A.F. Wals, 2004, Strategic generation with conjectured transmission price responses in a mixed transmission pricing system-Part 2: Application, *IEEE Transactions on Power Systems*, 19(2), 872-878, May.

Hobbs, B.F., C B. Metzler, J-S. Pang, 2000, Strategic gaming analysis for electric power systems: An MPEC approach, *IEEE Transactions on Power Systems*, 15(2), 638-644 May.

- Holmberg, P., 2008, Unique supply function equilibrium with capacity constraints, working paper, *Energy Economics*, 30(1), January 148-172.
- Holmberg, P., 2009, Numerical calculation of an asymmetric supply function equilibrium with capacity constraints, 199(1), November, 285-295.
- Holmberg, P. and D. Newbery, 2009, The supply function equilibrium and its policy implications for wholesale electricity auction, working paper 812, Research Institute of Industrial Economics.
- Holmberg, P., D. Newbery, and D. Ralph, 2010, Supply function equilibria: Step functions and continuous representations, working paper, Electricity Policy Research Group, University of Cambridge.
- Johnson, D.S., C.R. Aragon, L.A. McGeoch and C. Schevon, 1989, Optimization by Simulated Annealing: An Experimental Evaluation; Part I, Graph Partitioning, *Operations Research*, 37(6), 865–892.
- Kimbrough, S.O. and F.H. Murphy, 2009, Learning to Collude Tacitly on Production Levels by Oligopolistic Agents, *Computational Economics* 2009, 33(1), 47-78, February.
- Klemperer, P.D. and M. A. Meyer, 1989, Supply function equilibria in oligopoly under uncertainty, *Econometrica*, 57(6), 1243-1277, November.
- Kreps, D.M. and J.A. Scheinkman, 1983, Quantitative precommitment and Bertrand competition yield Cournot outcomes, *The Bell Journal of Economics*, 14(2) 326-337.
- Kruse, J.B., S. Rassenti, S.S. Reynolds; V.L. Smith, 1994, Bertrand-Edgeworth Competition in Experimental Markets, *Econometrica*, 62(2), 343-371, March.
- Krishna, V. and V.C. Ramesh, 1998, Intelligent agents for negotiations in market games, part 1: model, *IEEE Transactions on Power Systems*, 13(3), 1103-1108, August.
- Krishna, V. and V.C. Ramesh, 1998, Intelligent agents for negotiations in market games, part 2: application, *IEEE Transactions on Power Systems*, 13(3), 1109-1114, August.
- Kuenne, R.E., 1998, Price and nonprice rivalry in oligopoly: the integrated battleground, St. Martin's Press, New York, NY.
- Lave, L., S. Talukdar, K. Lye, and E. Subrahmanian, 2004. Designing electricity markets: Are freshmen or wind tunnels more useful? www.aeaweb.org/annual_mtg_papers/2005/0108_1015_1104.pdf.
- Liski, M. and J-P. Montero, 2006, Forward Trading and Collusion in Oligopoly, *Journal of Economic Theory*, 131(1), 212-230, November.
- Murphy, F.H. and Y. Smeers, 2005, Generation capacity expansion in imperfectly competitive restructured electricity markets, *Operations Research*, 53(4), 646-661, July August.
- Newbery, D., 2005, Electricity liberalisation in Britain: The quest for a satisfactory wholesale market design, *The Energy Journal*. p. 43-71.
- Rothkopf, M.H. 1999, Daily repetition: A neglected factor in the analysis of electricity auctions, *Electricity Journal*, 12(3) 61-79, April.
- Rudkevich, A., 2003, On the Supply Function Equilibrium and its Applications in Electricity Markets, working paper, Tabors Caramanis and Associates, Cambridge, MA.

- Sun, J. and L. Tesfatsion, 2006, Dynamic testing of wholesale power markets designs: an open-source agent-based framework, *Computational Economics*, 30(3) 291-327.
- Tirole, J., 1988, *The Theory of Industrial Organization*, MIT Press, Cambridge, MA.
- Ventosa, M., Á. Baíllo, A. Ramos, and M. Rivier, 2005. Electricity market modeling trends, *Energy Policy*, 33(7) 897-913 May.
- Visudhiphan, P., 2003, An agent-based approach to modeling electricity spot markets, dissertation Massachusetts Institute of Technology.
- von der Fehr, N-H. M. and D. Harbord, 1993, Spot market competition in the UK electricity industry, *The Economic Journal*, 103 531-546, May.
- Weidlich, A., 2008, Engineering interrelated electricity markets: An agent-based approach, Physical-Verlag, Heidelberg.
- Wilson, R., 2008, Supply function equilibrium in a constrained transmission system, *Operations Research*, 56(2), 369-382, March-April.
- Young, D., 2009, Kirchoffs Laws in Electricity Markets: Can Experimental Subjects Imitate Generators? University of Auckland Energy Centre.

A Probe and Adjust in Pseudo-Code

1. Set parameters $\delta, \varepsilon, \text{currentQuantity}, \text{epochLength}$
(Typically, $\varepsilon < \delta \ll \text{currentQuantity}$ and $\text{epochLength} \approx 50$.)
2. $\text{episodeCounter} \leftarrow 0$
3. $\text{returnsUp} \leftarrow []$ (Initialize returnsUp to an empty list.)
4. $\text{returnsDown} \leftarrow []$ (Initialize returnsDown to an empty list.)
5. Do forever:
6. $\text{episodeCounter} \leftarrow \text{episodeCounter} + 1$
7. $\text{bidQuantity} \sim U[\text{currentQuantity} - \delta, \text{currentQuantity} + \delta]$
(The agent's bidQuantity is drawn from the uniform distribution within the range $\text{currentQuantity} \pm \delta$.)
8. $\text{return} \leftarrow \text{Return-of } \text{bidQuantity}$
(The agent receives return from bidding bidQuantity .)
9. If ($\text{bidQuantity} \geq \text{currentQuantity}$) then:
 $\text{returnsUp} \leftarrow \text{Append } \text{return} \text{ to } \text{returnsUp}$
else:
 $\text{returnsDown} \leftarrow \text{Append } \text{return} \text{ to } \text{returnsDown}$
10. If ($\text{episodeCounter} \bmod \text{epochLength} = 0$) then:
(Epoch is over. Adjust episodeCounter and reset accumulators.)
 - (a) If ($\text{mean-of } \text{returnsUp} \geq \text{mean-of } \text{returnsDown}$) then:
 $\text{currentQuantity} \leftarrow \text{currentQuantity} + \varepsilon$
else:
 $\text{currentQuantity} \leftarrow \text{currentQuantity} - \varepsilon$
 - (b) $\text{returnsUp} \leftarrow []$
 - (c) $\text{returnsDown} \leftarrow []$
11. Loop back to step 5.

Figure 19: Pseudo-code for basic PROBE AND ADJUST

B Proofs

B.1 Lemma 1

Proof: We begin by analyzing the trials where Q_i^k is increased by δ . With j indexing the experiments in epoch k , let $\gamma_{-i}^{k,j} = \pm\delta$ be the random quantity adjustments from player $-i$'s experiments. The trial quantities in episode j are $Q_{-i}^{k,j} = Q_{-i}^k + \delta$ or $Q_{-i}^{k,j} = Q_{-i}^k - \delta$, depending on the outcome of the random choices by player $-i$. The profit for player i given an increase of δ is

$$[a - b(Q_i^k + Q_{-i}^k + \delta + \gamma_{-i}^{k,j}) - c_i^l](Q_i^k + \delta). \quad (12)$$

Rearranging terms, (12) becomes

$$[a - b(Q_i^k + Q_{-i}^k) - c_i^l]Q_i^k + [a - b(2Q_i^k + Q_{-i}^k) - c] \delta - b\delta^2 - b\gamma_{-i}^{k,j}(Q_i^k + \delta). \quad (13)$$

The change in profitability from the solution, (Q_1^k, Q_2^k) , at the start of epoch k is

$$[a - b(2Q_i^k + Q_{-i}^k) - c_i^l] \delta - b\delta^2 - b\gamma_{-i}^{k,j}(Q_i^k + \delta). \quad (14)$$

Let J_δ be the number of trials in the epoch with an increase of δ . Averaging over these trials in epoch k , we get

$$\sum_{j=1}^{J_\delta} \{[a - b(2Q_i^k + Q_{-i}^k) - c_i^l] \delta - b\delta^2 - b\gamma_{-i}^{k,j}(Q_i^k + \delta)\} / J_\delta = [a - b(2Q_i^k + Q_{-i}^k) - c_i^l] \delta - b\delta^2 - \sum_{j=1}^{J_\delta} [b\gamma_{-i}^{k,j}(Q_i^k + \delta)] / J_\delta. \quad (15)$$

Dividing through by δ results in

$$[a - b(2Q_i^k + Q_{-i}^k) - c_i^l] - b\delta - \frac{1}{\delta} \sum_{j=1}^{J_\delta} [b\gamma_{-i}^{k,j}(Q_i^k + \delta)] / J_\delta. \quad (16)$$

Note that the γ 's are iid Bernoulli trials taking on the values of δ and $-\delta$ with equal probability. Note also that with probability 1 as $J \rightarrow \infty$, $J_\delta \rightarrow \infty$. Thus, the sum, as $J_\delta \rightarrow \infty$ in the last term goes to 0. By assumption, $[a - b(2Q_i^k + Q_{-i}^k) - c_i^l] - b\delta > 0$. Consequently, with the increase to $Q_i^k + \delta$, in the limit the average profit from the sampling is higher than the profit at Q_i^k . So PROBE AND ADJUST will adjust up. The same logic shows that the average profit for all of the trials with a change of $-\delta$ in Q_i^k decreases relative to the profit at Q_i^k . Thus, at the end of epoch k we increase Q_i^k by ϵ . \square

B.2 Proof of Lemma 2

Proof: Initially, assume

$$0 < a - b(2Q_1^k + Q_2^k) - c_1^l < 3b\epsilon. \quad (17)$$

By Lemma 1 this implies an increase of ϵ for Q_1^k at the end of each such epoch. The maximum reduction in the left portion of the above expression occurs when player 2 also increases by ϵ and we get the expression

$$a - b[2(Q_1^k + \epsilon) + (Q_2^k + \epsilon)] - c_1^l. \quad (18)$$

By (17) we have

$$-3b\epsilon < a - b[2(Q_1^k + \epsilon) + Q_2^k + \epsilon] - c_1^l < 0. \quad (19)$$

From Lemma 1 we see that at iteration $k + 1$, $Q_1^{k+2} = Q_1^{k+1} - \epsilon$. Similarly, if the inequalities are reversed in (17), then $Q_1^{k+1} = Q_1^k - \epsilon$ and $Q_1^{k+2} = Q_1^{k+1} + \epsilon$.

Note that if the direction of change of Q_2^k , is different from Q_1^k , the sign of (18) does not necessarily change, but the inequality clearly holds and the lemma holds. \square

B.3 Proof of Theorem 1

Proof: From Lemma 1, we know that if at a trial solution (1) holds, then Q_i^k changes monotonically until (3) is satisfied for $i = 1, 2$. Clearly, at some iteration (3) is satisfied, otherwise, Q_i^k monotonically goes to 0 or infinity. Once (3) is satisfied, it is satisfied in all subsequent iterations. Restating (3), we have

$$-3b\epsilon < a - b(2Q_1^k + Q_2^k) - c_1^l < 3b\epsilon \quad (20)$$

and

$$-3b\epsilon < a - b(Q_1^k + 2Q_2^k) - c_2^l < 3b\epsilon. \quad (21)$$

Subtracting the definition of Q_i^* from the above equations we get

$$-3\epsilon < 2(Q_1^k - Q_1^*) + (Q_2^k - Q_2^*) < 3\epsilon \quad (22)$$

and

$$-3\epsilon < (Q_1^k - Q_1^*) + 2(Q_2^k - Q_2^*) < 3\epsilon. \quad (23)$$

Adding the above equations, we get

$$-2\epsilon < (Q_1^k - Q_1^*) + (Q_2^k - Q_2^*) < 2\epsilon \quad (24)$$

Subtracting (24) from (22) and (23), the result holds. \square

B.4 Proof of Corollary 1

Proof: Inequality (20) holds for all ϵ and the result follows. \square

B.5 Proof of Lemma 3

Proof: Without loss of generality assume $P_2^k > P_1^k$. In each episode we have four possible combinations of increases and decreases by δ . Whenever player 2 increases P_2^k by δ , $P_2^k + \delta > P_1^k \pm \delta$, it loses its whole second supply step, and its profits are $(P_1^k \pm \delta - c_2^1)q_2^1$. Whenever player 1 decreases P_1^k by δ , it wins the second supply step. Let $Q(P)$ be the quantity demanded at P . The player 1 profits are $(P_1^k \pm \delta - c_1^1)q_1^1 + (P_1^k \pm \delta - c_1^2)[Q(P_1^k \pm \delta) - q_1^1 - q_2^1]$. When player 1 increases its price and player 2 decreases its price, player 2 gains the supply step and its profits are $(P_2^k - \delta - c_2^1)q_2^1 + (P_2^k - \delta - c_1^2)[Q(P_2^k - \delta) - q_1^1 - q_2^1]$. Thus, the more profitable moves for both players are price cuts. Since both players cut prices the same amount,

$$|P_1^{k+1} - P_2^{k+1}| = |P_1^k - P_2^k| < 2\delta$$

and the result holds. \square

B.6 Proof of Lemma 4

Proof: As long as (5) holds, player 1 wins the market in every episode by having the lower price for all increases and decreases of δ and there is no randomness in the returns in the epoch. If $P_1^k < P_1^M$, profits are clearly higher with a price increase and vice versa. Thus, P_1^k moves monotonically to P_1^M . If (5) remains true when $|P_1^k - P_1^M| < 2\epsilon$, then P_1^k oscillates around P_1^M . Let this oscillation begin at epoch at k' . P_2^k moves in a random walk because the profits in each episode are $(P_1^k - c_2^1)q_2^1$ and are independent of the random choice $P_2^k \pm \delta$. Since this is a random walk without any drift, with probability 1 there is a $k > k'$ such that $P_1^M + b\epsilon > P_2^k$ and thus, $|P_1^k - P_2^k| < 2\delta$.² \square

B.7 Proof of Theorem 2

Proof: By Lemma 4, for some k , $|P_1^k - P_2^k| < 2\delta$. By Lemma 3, from iteration k on P_i^k decreases its price until $P_1^k < c_2^2$. From the proof of Lemma 4, the profits at $c_2^2 > P_1^k + \delta$ are greater than the profits at $c_2^2 > P_1^k - \delta$ and the results hold. \square

B.8 Proof of Lemma 5

Proof: Since player 2 is on its vertical, Q_2 can be taken as fixed, so the price p set by player 1 is linked to the inverse demand curve as follows.

$$p = a - b(Q_1 + Q_2). \tag{25}$$

Thus,

$$\frac{a-p}{b} - Q_2 = Q_1 \tag{26}$$

²We assume δ is chosen close to $b\epsilon$. With δ too small relative to ϵ the trial solutions will likely just wander.

and the profit for player 1 is

$$\left[\frac{a-p}{b} - Q_2 \right] (p - c_1^2). \quad (27)$$

Taking the derivative with respect to p and setting it equal to 0, we get

$$0 = \frac{a}{b} - \frac{2p}{b} - Q_2 + \frac{c_1^2}{b} = a - 2p - bQ_2 + c_1^2. \quad (28)$$

Substituting (25) for p in (28), we have

$$a - 2a + 2b(Q_1 + Q_2) - bQ_2 + c_1^2 = 0 \quad (29)$$

or

$$a - 2bQ_1 - bQ_2 - c_1^2 = 0 \quad (30)$$

This is the Cournot condition for player 1. \square

B.9 Proof of Lemma 6

Proof: As in Lemma 1, we can use expected values as J gets large. We calculate the expected value in the limit of the difference in returns as follows.

$$(P_1^k + b\delta - c_2^1)(Q_2^k + \delta) - (P_1^k + b\delta - c_2^1)(Q_2^k - \delta) = (P_1^k + b\delta - c_2^1)2\delta. \quad (31)$$

Similarly,

$$(P_1^k - b\delta - c_2^1)(Q_2^k + \delta) - (P_1^k - b\delta - c_2^1)(Q_2^k - \delta) = (P_1^k - b\delta - c_2^1)2\delta. \quad (32)$$

Adding the equations together gives us the expected return.

$$2(P_1^k - c_2^1)2\delta > 0, \quad (33)$$

Since $Q_2^k < q_2^1$, Q_2^k can increase and the result holds. \square

B.10 Proof of Lemma 7

Proof: Initially, assume that

$$0 < a - 2bQ_1^k - bQ_2^k - c_1^2 < 3b\epsilon. \quad (34)$$

Then by Lemma 6 player 1 lowers the price by $b\epsilon$. This increases the total quantity demanded by ϵ . The maximum change in the left hand portion of (34) happens when $Q_1^{k+1} = Q_1^k + 2\epsilon$ and $Q_2^{k+1} = Q_2^k + \epsilon$. When that happens we have

$$a - 2bQ_1^{k+1} - bQ_2^{k+1} - c_1^2 - a + 2bQ_1^k + bQ_2^k + c_1^2 = -3b\epsilon. \quad (35)$$

which leads to

$$-3b\epsilon < a - 2bQ_1^{k+1} - bQ_2^{k+1} - c_1^2 < 0. \quad (36)$$

Note that with an increase in Q_2^k , the difference in (33) is $-\epsilon$. With the inequality reversed in (34) and using the same logic, we get (36) with the inequalities reversed through applying the same logic and the Lemma holds. \square

B.11 Proof of Lemma 8

Proof: Assume that (7) holds for $k = k'$. Then at iteration $k+1$, $p^{k+1} = p^k - b\epsilon$, $Q_1^{k+1} + Q_2^{k+1} = Q_1^k + Q_2^k + \epsilon$, and $Q_1^{k+1} \geq Q_1^k$. This means that

$$a - 2bQ_1^{k+1} - bQ_2^{k+1} - c_1^2 \leq a - 2bQ_1^k - bQ_2^k - c_1^2 - \epsilon < a - 2bQ_1^k - bQ_2^k - c_1^2. \quad (37)$$

Applying the same logic to (8), we have the inequalities in (37) reversed, and the result holds.

Note that when $p^{k+1} = p^k + b\epsilon$ and $p^{k+1} > c_2^1$,

$$Q_1^{k+1} + Q_2^{k+1} = Q_1^k + Q_2^k - \epsilon. \quad (38)$$

Since $Q_2^{k+1} = Q_2^k + \epsilon$,

$$Q_1^{k+1} = Q_1^k - 2\epsilon. \quad (39)$$

Also, when $p^{k+1} = p^k - b\epsilon$ and $p^{k+1} > c_2^1$,

$$Q_1^{k+1} + Q_2^{k+1} = Q_1^k + Q_2^k + \epsilon. \quad (40)$$

Since $Q_2^{k+1} = Q_2^k + \epsilon$,

$$Q_2^{k+1} = Q_2^k. \square \quad (41)$$

B.12 Proof of Lemma 9

Proof: By Lemma 6, Q_2^k increases by ϵ when $q_1^1 + \epsilon < Q_1^k < q_1^2 - \epsilon$. Since q_2^1 is finite, if $q_1^1 + \epsilon < Q_1^k < q_1^2 - \epsilon$ for all trials between k and k' , then $Q_2^{k'} \geq q_2^1 - \epsilon$ for some k' , and (i) holds. If $Q_2^{k'} < q_2^1 - \epsilon$ is true for iterations k through k' , then either $Q_1^{k'} \leq q_1^1 + \epsilon$ is true or $Q_1^{k'} \geq q_1^2 - \epsilon$. We show that the latter case is not possible. From the proof of 8, we know that $Q_1^{k+1} + Q_2^{k+1} \leq Q_1^k + Q_2^k + \epsilon$. Since $Q_2^{k+1} = Q_2^k + \epsilon$, it follows that $Q_1^{k+1} \leq Q_1^k$. Thus, (ii) is the only other possibility at k' . \square

B.13 Proof of Theorem 3

Proof: The step-function marginal-cost curves can be used in a Cournot game to find the Cournot equilibrium. This equilibrium exists and is unique using the standard theory since the cost curves are monotonically nondecreasing and the optimization problem for each player in a Cournot game is convex. Note that if the quantity offered is on the vertical of the marginal cost curve, the right and left derivatives of the profit function at this point exist and are of the right sign. We now show that this equilibrium is a local equilibrium in the supply-function game.

Let c_i^c be the marginal cost for player i at Q_i^c . By the Cournot condition, $P^c > c_i^c$. If each player defines its offer curve by setting the offer prices equal to marginal cost for $q_i \leq Q_i^c$ and some large number R for $q_i > Q_i^c$, then each of these offer curves is a best response to the other offer curve when prices and quantities are perturbed a small amount, i.e., are in the set of local offer curves. \square

B.14 Proof of Lemma 10

Proof: Since $Q_{-i}^b = q_{-i}^1$ for $P^b = c_{-i}^2$, player $-i$'s quantity is fixed and player i sees the Cournot equilibrium condition on its profits. If

$$a - 2bQ_i^b - Q_{-i}b - c_i^2 > 0, \tag{42}$$

then player i can increase its profits by lowering P^b and increasing Q_i^b and the Bertrand solution is not a local equilibrium. The condition (10) follows from noting that the limiting price is $c_{-i}^2 = a - bQ_i^b - bQ_{-i}^b$. Thus, the result holds. \square

C Computational Model

Our computational model is embodied in the NetLogo program `SupplyCurveBidding.nlogo`, which may be accessed at <http://asadai.wharton.upenn.edu/nlogo/SupplyCurveBidding.html>. We distinguish two kinds of parameters in our model: environment parameters and agent parameters. Tables 1 and 2 list the default values for these parameters. We call these the Base Case Parameterization. Settings for specific runs can be identified as deviations from these default values.

The core structure of the computational model may be seen by reviewing the parameters and what they do, along with the program’s user interface, Figure 20. We begin with the environment parameters. Demand is an exogenously-given linear function, with price as a function of quantity. The demand curve is specified by the quantity intercept (the quantity demanded at 0 price, `QIntercept`) and the (downward) slope of the line (`DSlope`). `EpisodesPerRun` is the number of rounds of play in a single run under the settings. `NumberOfReplications` is the number of runs undertaken with the same parameter settings. `RunningAverageLength` is the number of periods of data collected to compute the various running average statistics. When `Logging?` is set to true, run data are written to the file `runsOutput.txt` in CSV (comma separated values) format with headers for easy loading into spreadsheets and statistical programs. (We used R.)

<code>QIntercept</code> (Quantity intercept of demand function)	200
<code>DSlope</code> (Negative of the slope of the demand function)	1.0
<code>EpisodesPerRun</code> (Number of episodes in a run)	50,000
<code>RunningAverageLength</code> (Running average length)	2,000
<code>RandomNumberSeed</code> (Random number seeded with)	System Clock
<code>NumberOfReplications</code> (Number of runs)	30
<code>Logging?</code> (Are the results to be logged?)	True

Table 1: Default settings of the environment parameters, Base Case Parameterization.

NumberOfFirms	2
Cost1Firm1	4
Cost1Firm2	4
MaxQ1Firm1	20
MaxQ1Firm2	20
Cost2Firm1	15
Cost2Firm2	15
MaxQ2Firm1	100
MaxQ2Firm2	100
EpochLengthFirm1	100
EpochLengthFirm2	100
InitialBasePrice1Firm1	16
InitialBasePrice1Firm2	16
InitialBasePrice2Firm1	150
InitialBasePrice2Firm2	150
InitialBasePrice3Firm1	150
InitialBasePrice3Firm2	150
QuantityBreak1Firm1	0.5
QuantityBreak1Firm2	0.5
QuantityBreak2Firm1	0.9
QuantityBreak2Firm2	0.9
UpdateTypeFirm1	Own Returns
UpdateTypeFirm2	Own Returns
Patience	1.05
QuantityDeltaFirm1	0.01
QuantityDeltaFirm2	0.01
QuantityEpsilonFirm1	0.01
QuantityEpsilonFirm2	0.01
PriceDeltaFirm1	0.2
PriceDeltaFirm2	0.2
PriceEpsilonFirm1	0.2
PriceEpsilonFirm2	0.2

Table 2: Default settings of the agent parameters, Base Case Parameterization.

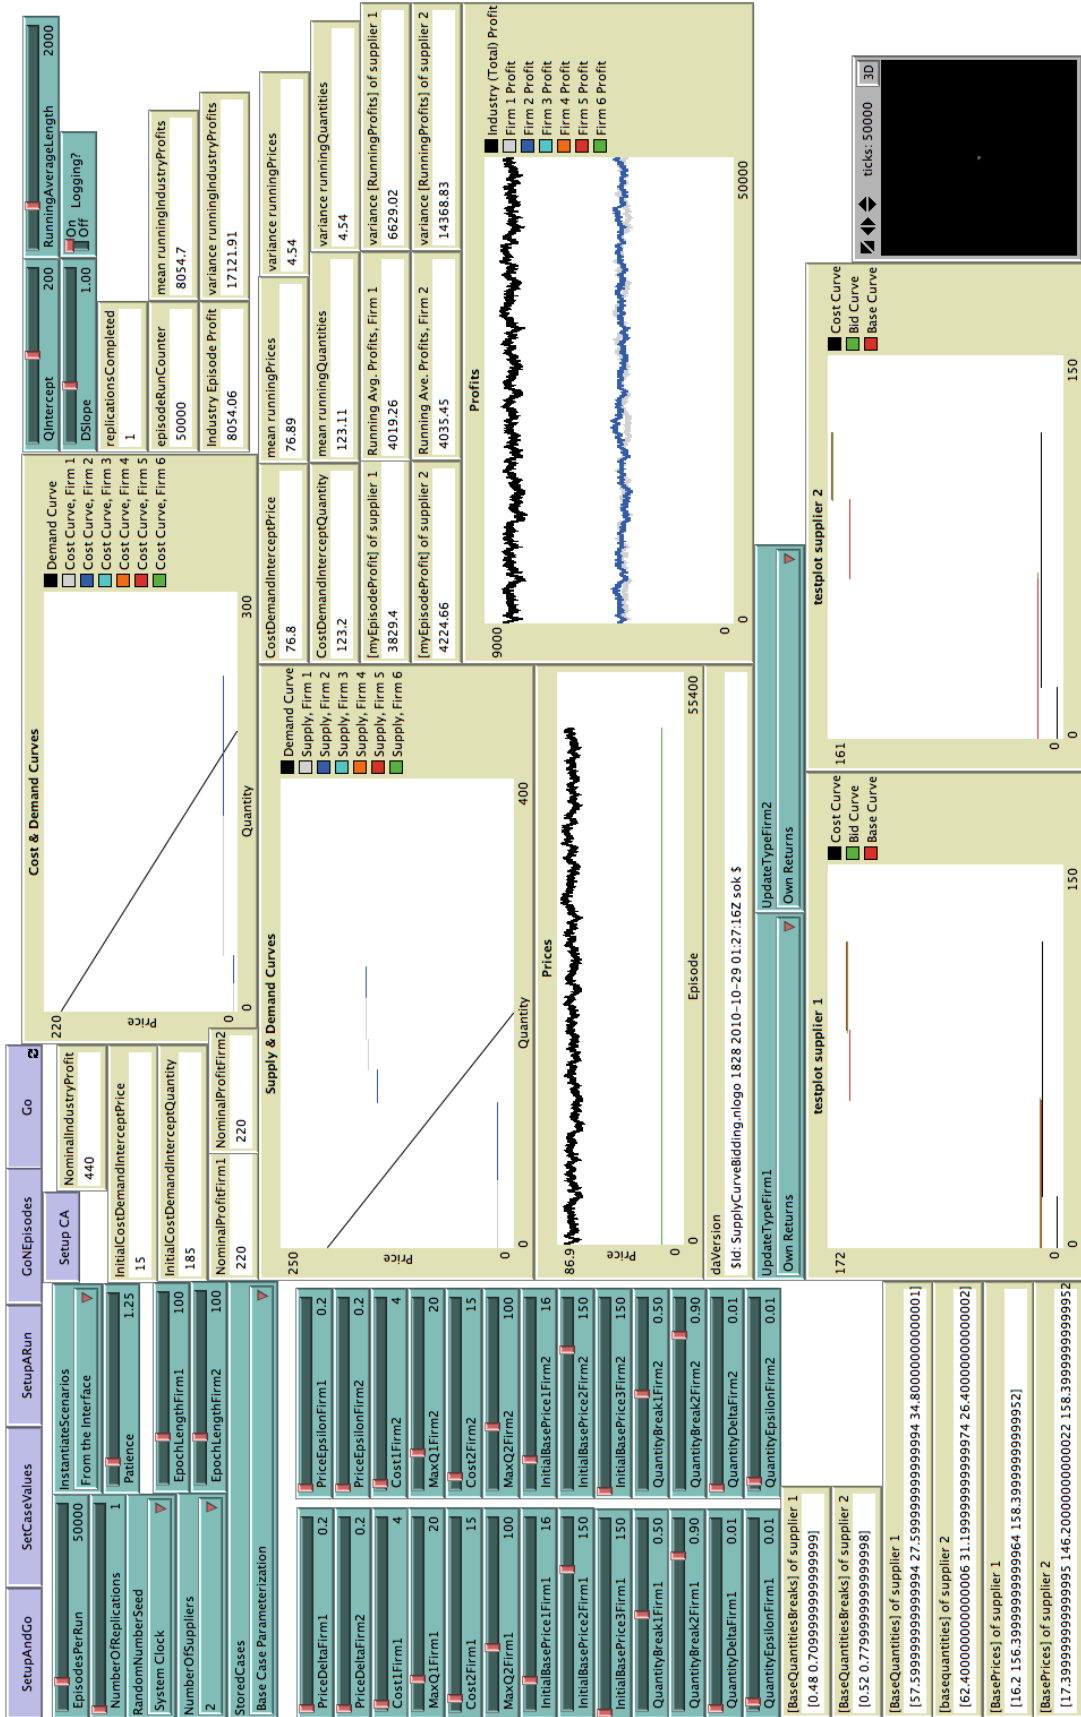


Figure 20: Base Case Parameterization, representative run (both suppliers update with Own Returns)

The agent parameters in Table 2 and the program interface in Figure 20 may be best understood in the context of a narrative of how the model works.

The program allows any number of agents (supplier firms), each of which may have one or more supply plants. Agents may bid one or more (price, quantity) pairs, committing thereby to supplying up to the stated quantity at the stated price per unit. This is subject to the constraint that the total quantity bid must equal the total quantity available from the agent's plants. Additionally, bids are constrained to be at or above cost for all of the capacity required to fulfill them. An agent making bids with N (price, quantity) pairs will use N base prices and $N - 1$ boundaries between them for specifying quantities. We call these latter values *quantity break points*. There is thus a total of $2N - 1$ parameters which are learned via PROBE AND ADJUST for each supplier.

In our Base Case Parameterization, for example, there are two agents, firm 1 and firm 2. Each agent has two plants with distinct cost and quantity characteristics. For firm 1, $\text{Cost1Firm1} = 4$ and $\text{Cost2Firm1} = 15$, meaning that the two plants for firm 1 have associated costs of 4 and 15. Their capacities are 20 and 100 respectively. Internally, the program represents the costs and capacities of the plants in lists (**costsofplants** and **capacitiesofplants**), so that any number of plants may be modeled. The agents in the Base Case each make bids consisting of three (price, quantity) pairs. These are generated by three price values and two quantity breakpoint values. In the Base Case (Table 2 and Figure 20) three initial base prices for firm 1 are 16, 150, and 150. The initial quantity breakpoints are at 0.5 and 0.9. If these values are bid as is (without perturbation by PROBE AND ADJUST), then, since the total capacity is 120, the bid would consist of the three price/quantity pairs: (16, 60), (150, 48), and (150, 12). Internally, an agent's base prices are stored in a list, **BasePrices**, as are its base quantities breaks, **BaseQuantitiesBreaks**. Consequently, an agent may be initialized to bid an arbitrary number of (price, quantity) pairs.

An episode (round of play) begins with each agent consulting its N **BasePrices** and $(N-1)$ **BaseQuantitiesBreaks** and using PROBE AND ADJUST to create its current bid as N (price, quantity) pairs. The market agent sums the bids from each firm to create the supply curve for the episode. The market agent then clears the market by determining the intersection of the supply and demand curves, and thereby the prevailing price and quantities. Bidders are rewarded according to the prevailing price and the quantity taken by the market. Upon receiving their rewards from the market, each agent collects information and prepares for the next episode. The **EpochLength** parameter for each agent determines how often the base quantities and prices are updated. In the Base Case, the firms both update every 100 episodes.

Strategically, the most interesting parameters or decisions for an agent are its update type (see **UpdateTypeFirm1**). The two most important update types are Own Returns and Market Returns Constrained by Own Returns. An Own Returns player updates the adjustable quantities in its bids (i.e., the values in **BasePrices** and **BaseQuantitiesBreaks**) with PROBE AND ADJUST, using its profits (own returns). That is, an Own Returns supplier keeps track of its returns from its bids on each of its $(2N - 1)$ bid parameters every episode. Acting on this performance measure is a form of non-cooperative behavior where the decisions are based on

the outcomes of a learning process rather than an explicit optimization. The return in a episode is coded high or low for each of the $(2N - 1)$ bid parameters, depending upon whether the bid value of the parameter is above or below its base value for the episode. Learning and adjustment occurs for each price and quantity on the offer curve separately.

A Market Returns player instead updates based on maximizing the industry's returns (net of costs). This is an altruistic strategy and obviously is exploitable, but it does lead to superior returns for the industry. The Market Returns Constrained by Own Return (MR-COR) update type uses Market Returns to update its adjustable quantities provided its profits are not too much below the average profits of the industry. Here is the relevant code fragment.

```
((MyMean myEpisodeProfits) * Patience) >=
      (MyMean industryEpisodeProfits) / count suppliers)
```

If it evaluates to true, a MR-COR player will update altruistically, based on Market Returns; if it evaluates to false, a MR-COR player will update based on Own Returns. (Note: In the NetLogo code we use the term `supplier` for the breed (type) of agent that represents supplier firms. `count suppliers` is simply the number of firms in the market, 2 for the Base Case.) That is, if on average the player's profits (returns net of costs) are not much below the industry's profits in the (player's) epoch just concluded, then the player updates on the "Market Returns" principle; otherwise the player updates on the basis of the "Own Returns" principle. This update criterion is a form of cooperative behavior with punishments for a lack of cooperation based on past experience and without formal and illegal communication.

D Additional Computational Results

We present in this section various supplementary computational materials in support of findings described in the body of the paper. We have made the program available on the Web as an applet that can be run from any Java-enabled browser: <http://asadai.wharton.upenn.edu/nlogo/SupplyCurveBidding.html>. In addition, the program itself is freely available for downloading from the Web site. Users who download the program will need to install NetLogo in order to run the program. NetLogo is freely available at <http://ccl.northwestern.edu/netlogo/>. Our program has been tested on version 4.1.2 of NetLogo, which was current at the time of posting.

D.1 Run Summaries

Tables 3–7 present result summaries for three different experiments with the Base Case parameterization, undertaken by varying the update types of the two supplier firms. The abbreviations in these and subsequent tables have the following meanings. Mean RIP: the mean of the running industry profits at the end of the

run. With `RunningAverageLength` set to 2,000 in the Base Case (see Table 1), we record at the end of each run the mean industry profits for the last 2,000 episodes. Mean RIP is the mean of these values, averaged over the `NumberOfReplications` runs (here 30). Mean RP: mean of the running prices, computed similarly to Mean RIP. Mean RQ: mean of the running quantities. Mean RPSX: mean of the running profits for supplier X. Again, these are computed in a manner similar to Mean RIP.

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
Mean running industry profit (RIP)	7907	7982	8056	8046	8094	8261
Mean running price (RP)	74.60	75.73	76.92	76.78	77.55	80.45
Mean running quantity (RQ)	119.6	122.5	123.1	123.2	124.3	125.4
Mean running profit supplier 1 (RPS1)	3920	3998	4052	4045	4089	4181
Mean running profit supplier 2 (RPS2)	3891	3955	4011	4001	4030	4148

Table 3: Base Case Parameterization (both suppliers update with Own Returns). See Figure 20 in §C.

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
Mean running industry profit (RIP)	234.2	240.1	373.9	350.1	411.8	803.7
Mean running price (RP)	15.08	15.10	15.18	15.27	15.20	18.12
Mean running quantity (RQ)	181.9	184.8	184.8	184.7	184.9	184.9
Mean running profit supplier 1 (RPS1)	21.82	124.46	156.88	163.38	208.28	301.91
Mean running profit supplier 2 (RPS2)	109.8	122.0	155.8	186.7	213.9	781.9

Table 4: Bertrand250Both Parameterization (both suppliers update with Own Returns). See Figure 5.

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
Mean running industry profit (RIP)	2800	5386	5439	5357	5493	5652
Mean running price (RP)	28.80	47.41	47.86	47.28	48.31	49.67
Mean running quantity (RQ)	150.3	151.7	152.1	152.7	152.6	171.2
Mean running profit supplier 1 (RPS1)	1279	1280	1284	1317	1286	2292
Mean running profit supplier 2 (RPS2)	508	4099	4158	4041	4212	4369

Table 5: Q2Firm1Is250 Parameterization (both suppliers update with Own Returns). See Figure 32 in §D.

The Tables 9 and 10 summarize results from 30 replications each.

D.2 Screen Dumps

What follows is a series of “screen dumps” from runs of `SupplyCurveBidding.nlogo`. Captions for the figures contain the information required to duplicate the run displayed, albeit with a different random number

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
Mean running industry profit (RIP)	8994	8994	8995	8995	8995	8995
Mean running price (RP)	107.3	107.5	107.5	107.5	107.6	107.7
Mean running quantity (RQ)	92.35	92.40	92.45	92.46	92.50	92.73
Mean running profit supplier 1 (RPS1)	3614	4283	4494	4530	4936	5264
Mean running profit supplier 2 (RPS2)	3730	4059	4501	4464	4711	5380

Table 6: Base Case Parameterization (both suppliers update with MR-COR).

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
Mean running industry profit (RIP)	7968	8093	8136	8150	8221	8318
Mean running price (RP)	75.56	77.53	78.28	78.53	79.71	81.52
Mean running quantity (RQ)	118.5	120.3	121.7	121.5	122.5	124.4
Mean running profit supplier 1 (RPS1)	4200	4273	4301	4302	4337	4391
Mean running profit supplier 2 (RPS2)	3737	3814	3849	3848	3884	3941

Table 7: Base Case Parameterization (firm 1 updates with Own Returns, firm 2 with MR-COR).

stream.

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
Mean running industry profit (RIP)	5179	5324	5392	5406	5497	5648
Mean running price (RP)	45.72	46.90	47.47	47.59	48.35	49.64
Mean running quantity (RQ)	150.4	151.7	152.5	152.4	153.1	154.3
Mean running profit supplier 1 (RPS1)	3897	4037	4099	4106	4198	4366
Mean running profit supplier 2 (RPS2)	1282	1285	1286	1300	1290	1437

Table 8: P2BothIs25 Parameterization (firm 1 updates with Own Returns, firm 2 with MR-COR). See Figure 26 for a representative run.

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
Mean running industry profit (RIP)	5480	5636	5965	6167	6654	7271
Mean running price (RP)	48.20	49.53	52.44	54.60	59.10	65.97
Mean running quantity (RQ)	134.0	140.9	147.6	145.4	150.5	151.8
Mean running profit supplier 1 (RPS1)	1284	1398	1666	1950	2465	3695
Mean running profit supplier 2 (RPS2)	3482	4112	4225	4217	4387	4616

Table 9: Q2Firm1Is250 Parameterization (both suppliers update with MR-COR). See Figure 36 in §D.

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
Mean running industry profit (RIP)	5325	5676	6121	6117	6446	7361
Mean running price (RP)	46.91	49.88	53.89	54.06	57.00	67.07
Mean running quantity (RQ)	132.9	143.0	146.1	145.9	150.1	153.1
Mean running profit supplier 1 (RPS1)	1295	1426	1735	1774	1983	3041
Mean running profit supplier 2 (RPS2)	4007	4152	4327	4343	4482	4767

Table 10: Q2Firm1Is250 Parameterization (firm 1 updates with Own Returns, firm 2 with MR-COR). See Figure 37 in §D.

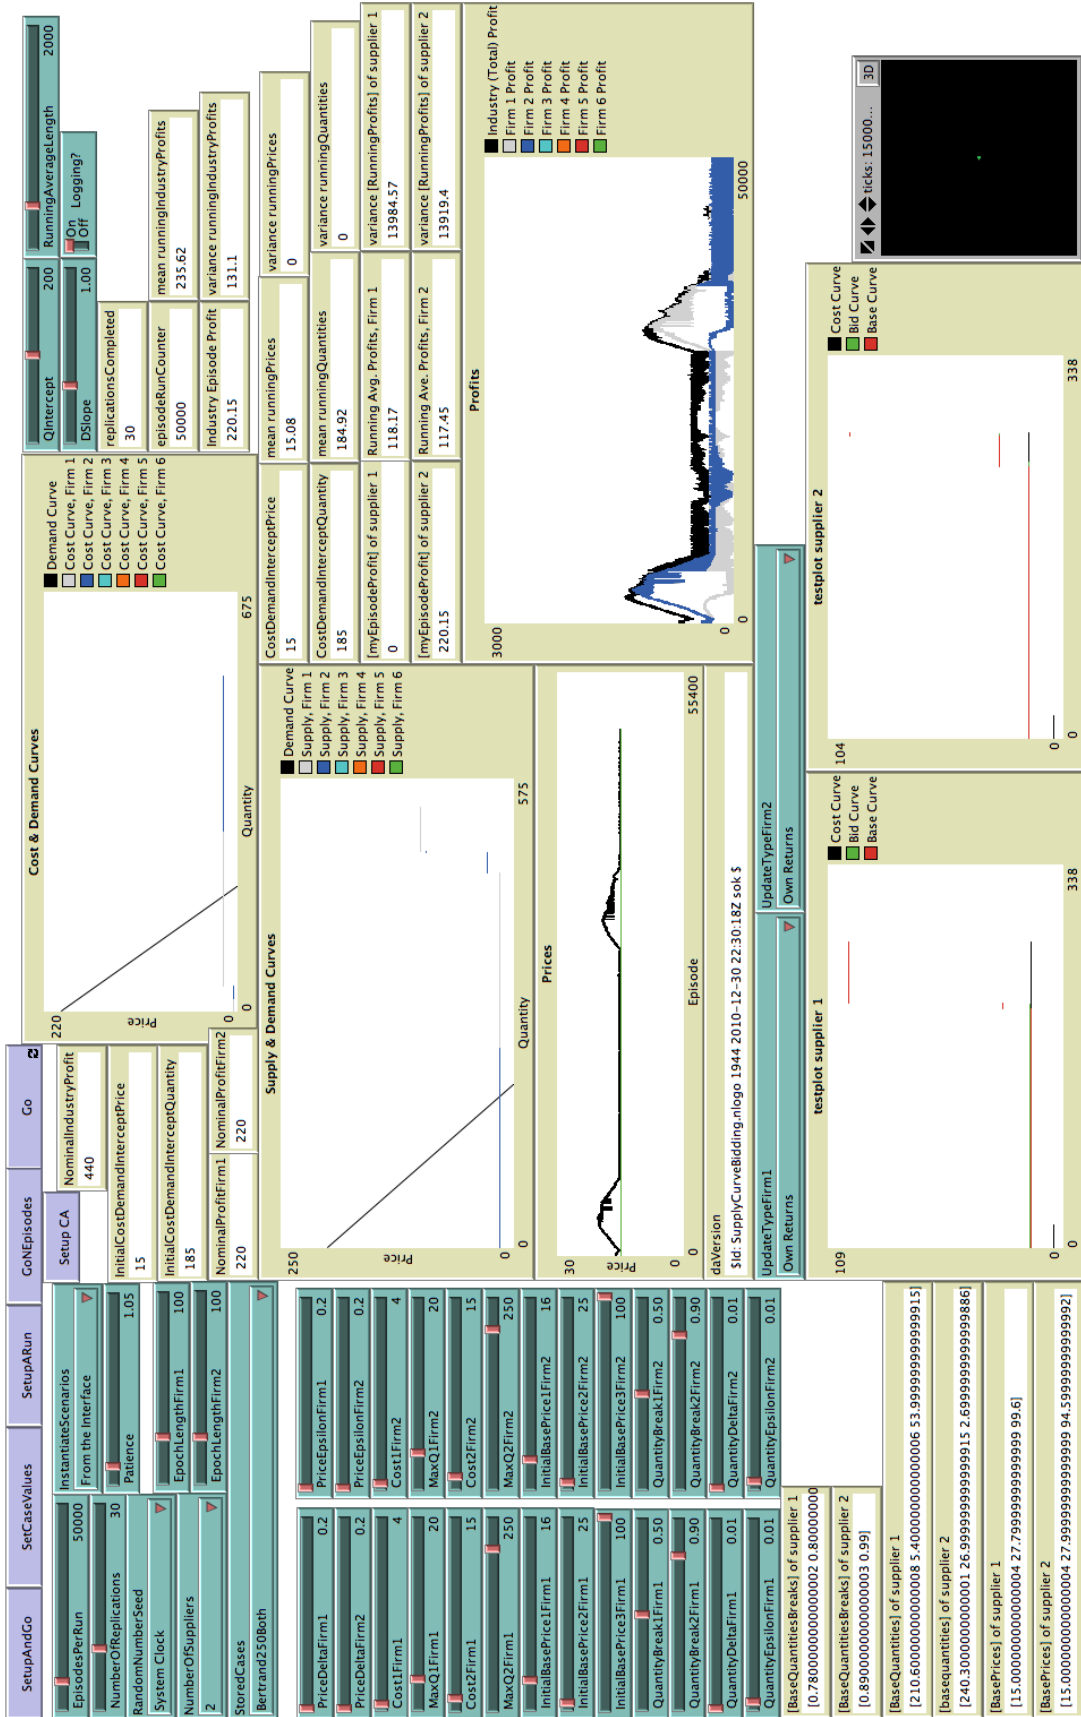


Figure 21: Bertrand250Both stored case, representative run (both suppliers update with Own Returns)

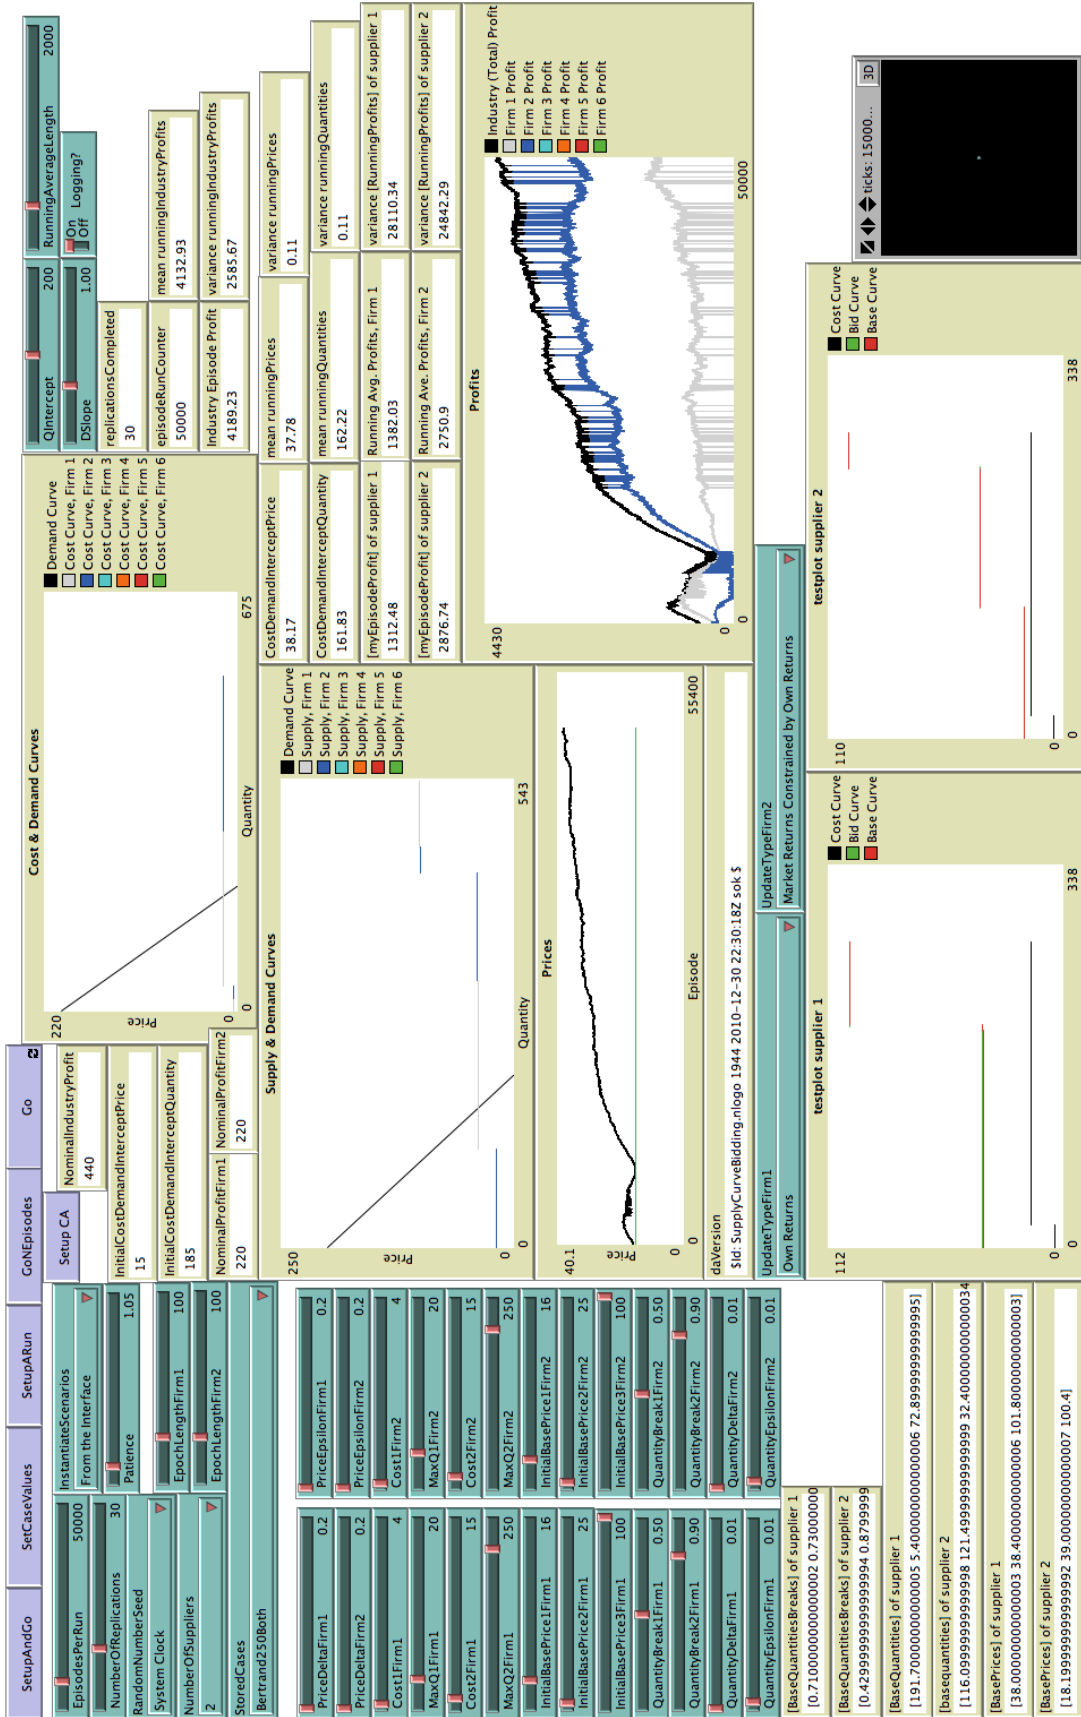


Figure 23: Bertrand250Both stored case, representative run (firm 1 updates with Own Returns, firm 2 with MR-COR)

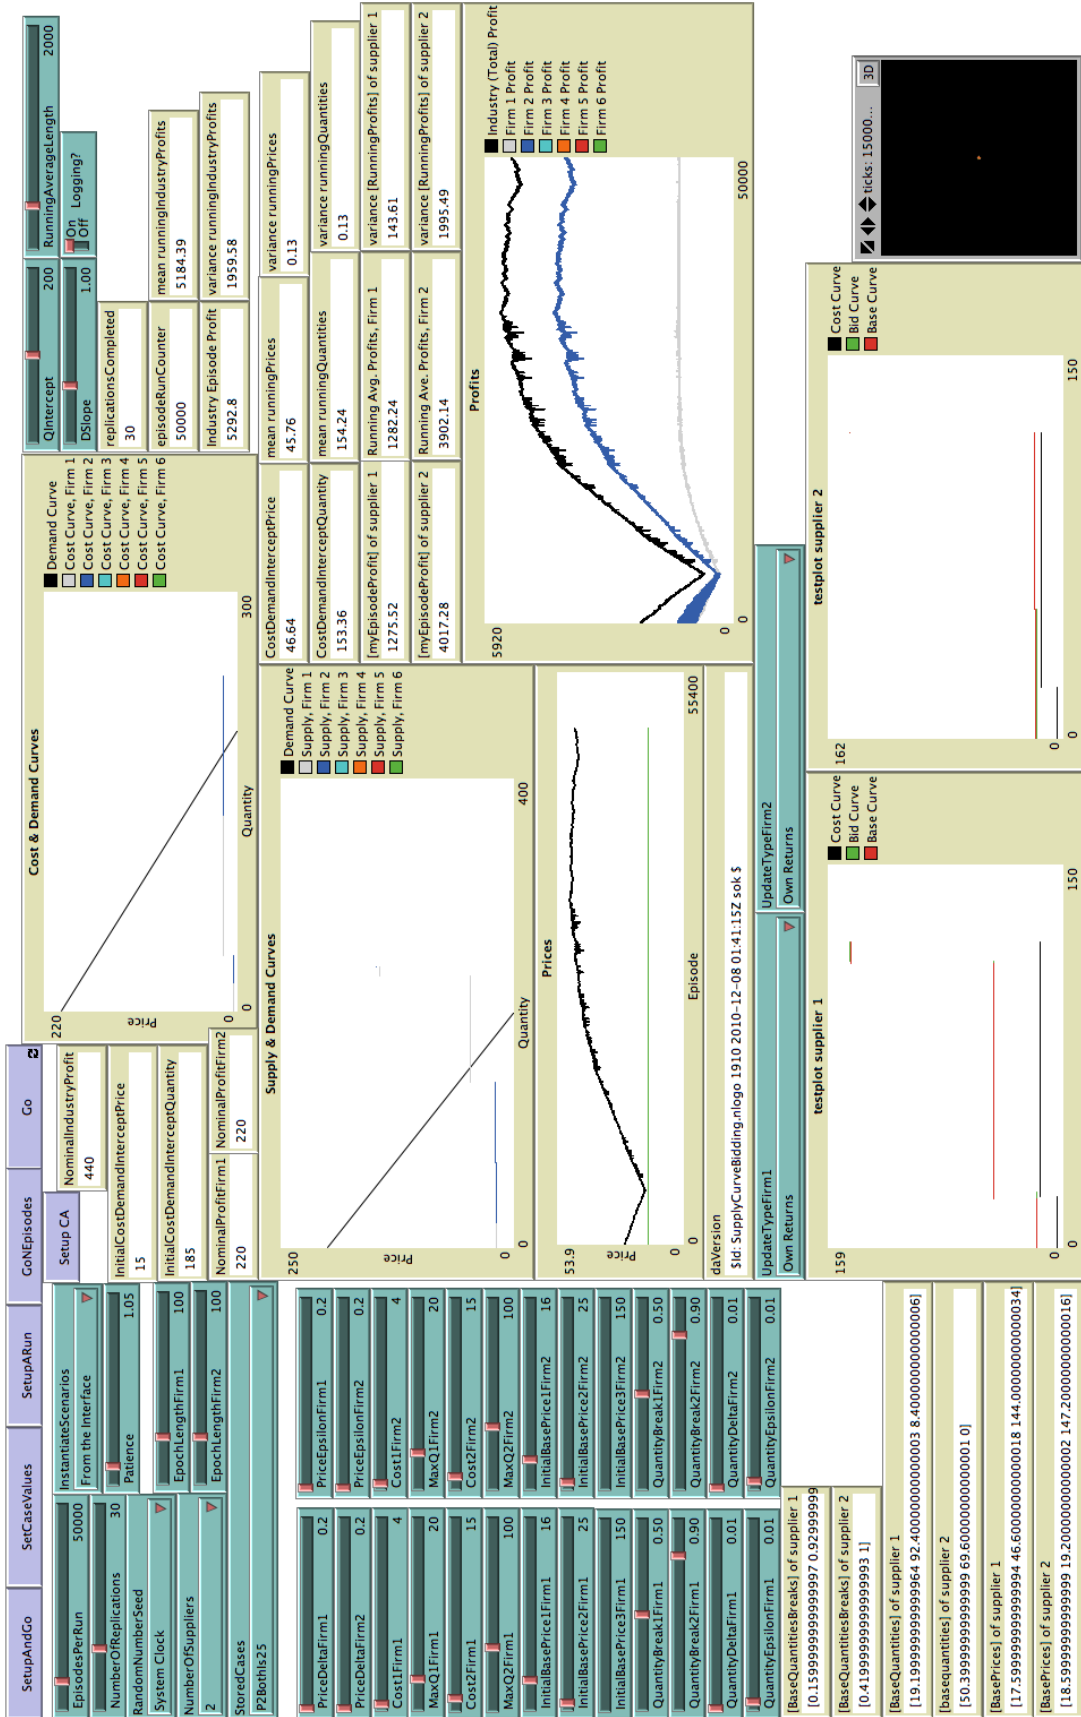


Figure 24: P2BothIs25 stored case, representative run (both suppliers update with Own Returns)

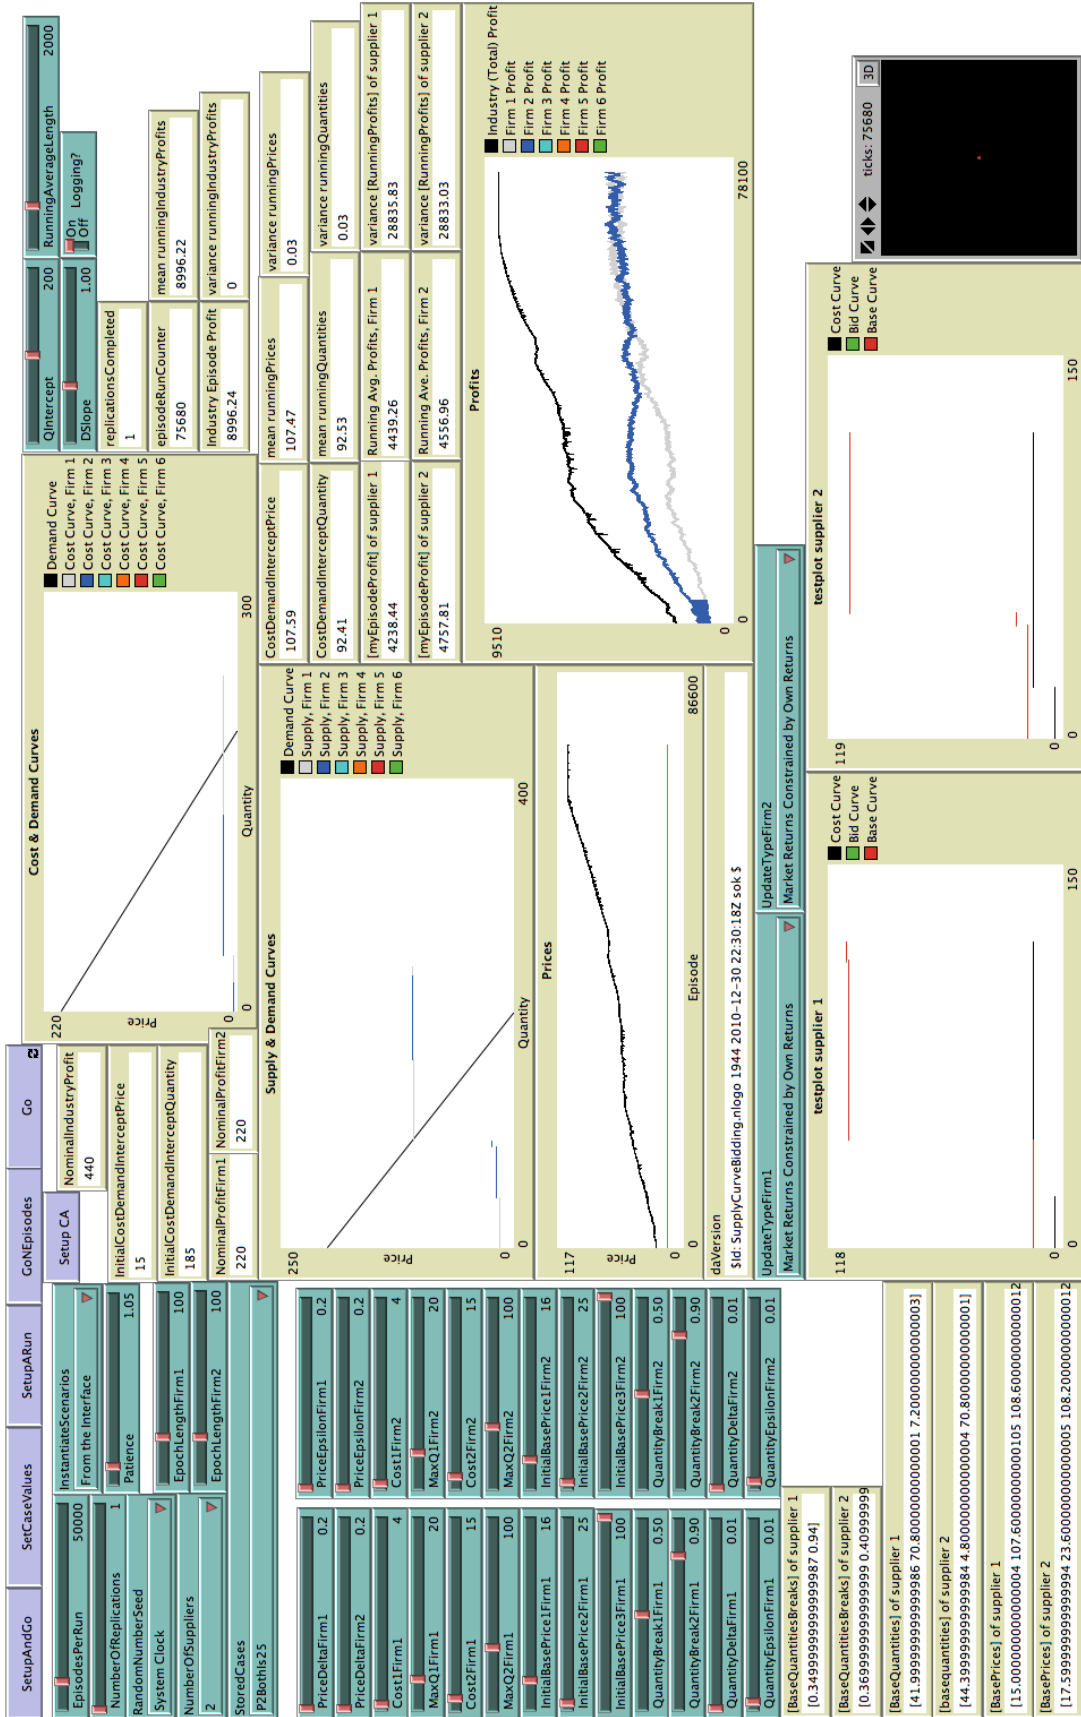


Figure 25: P2BotHls25 stored case, representative run (both suppliers update with MR-COR)

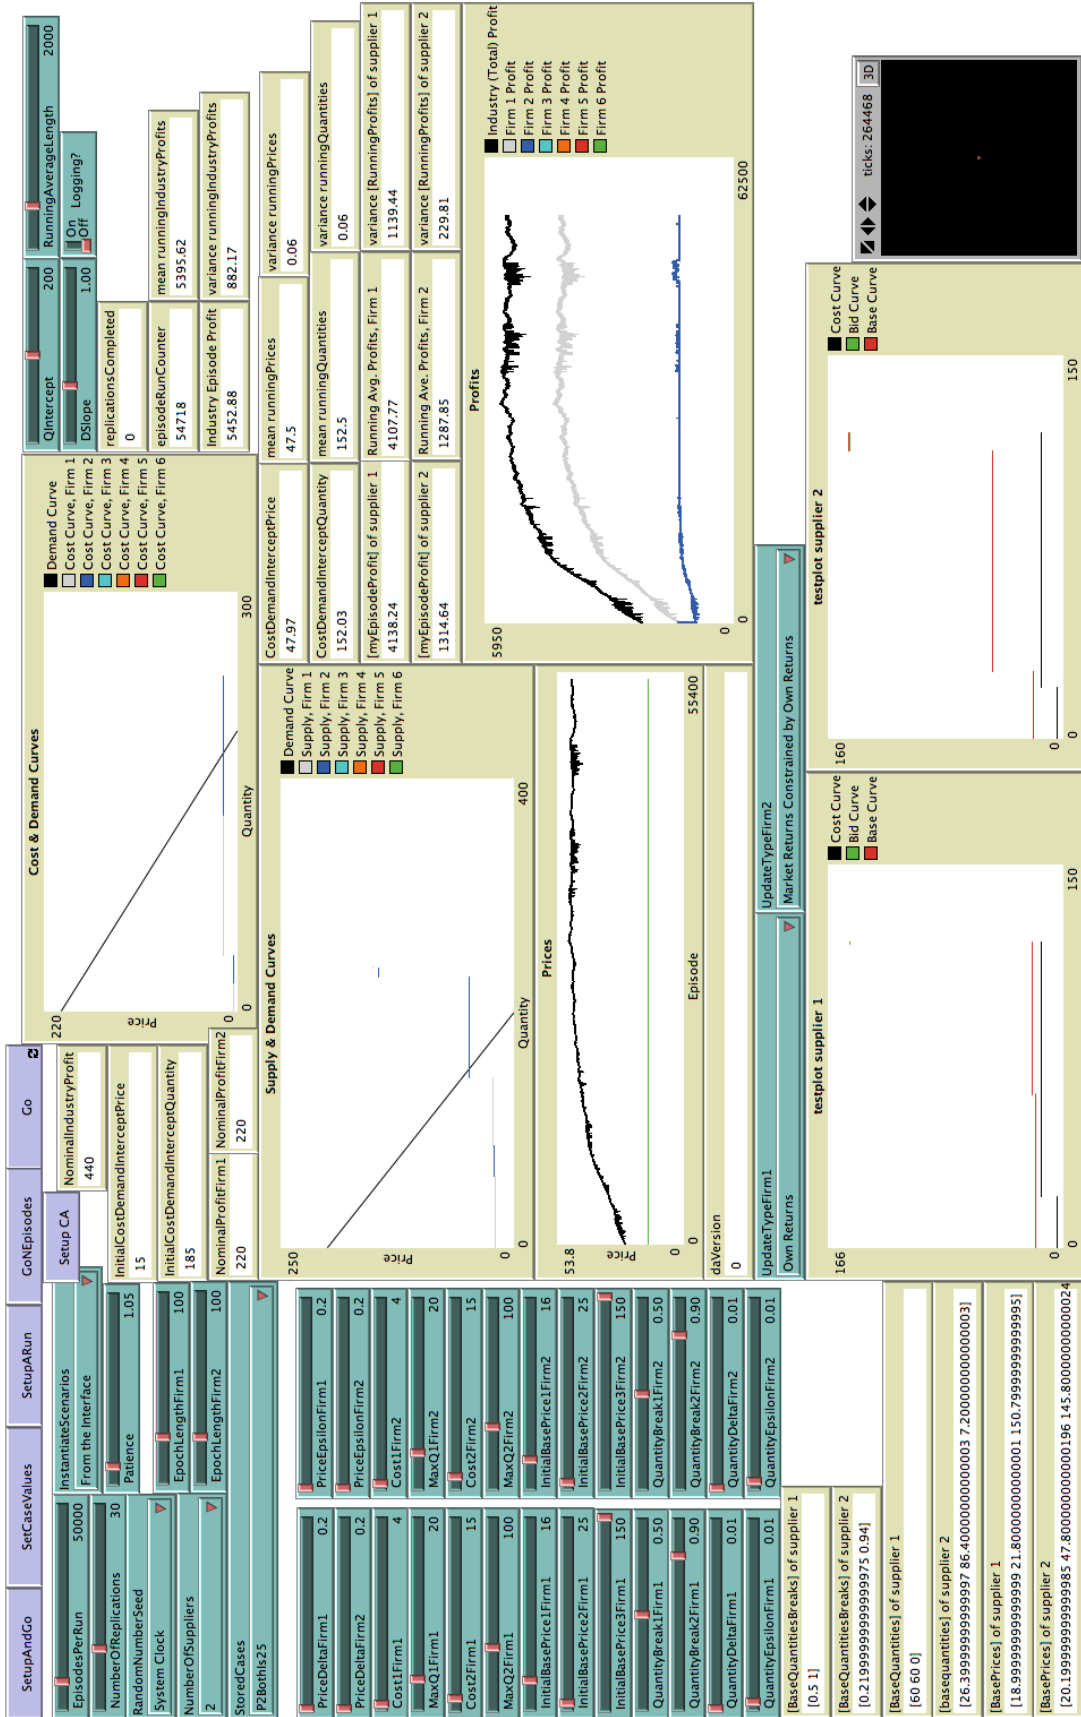


Figure 26: P2BothIs25 stored case, representative run (firm1 updates with Own Returns, firm 2 with MR-COR)

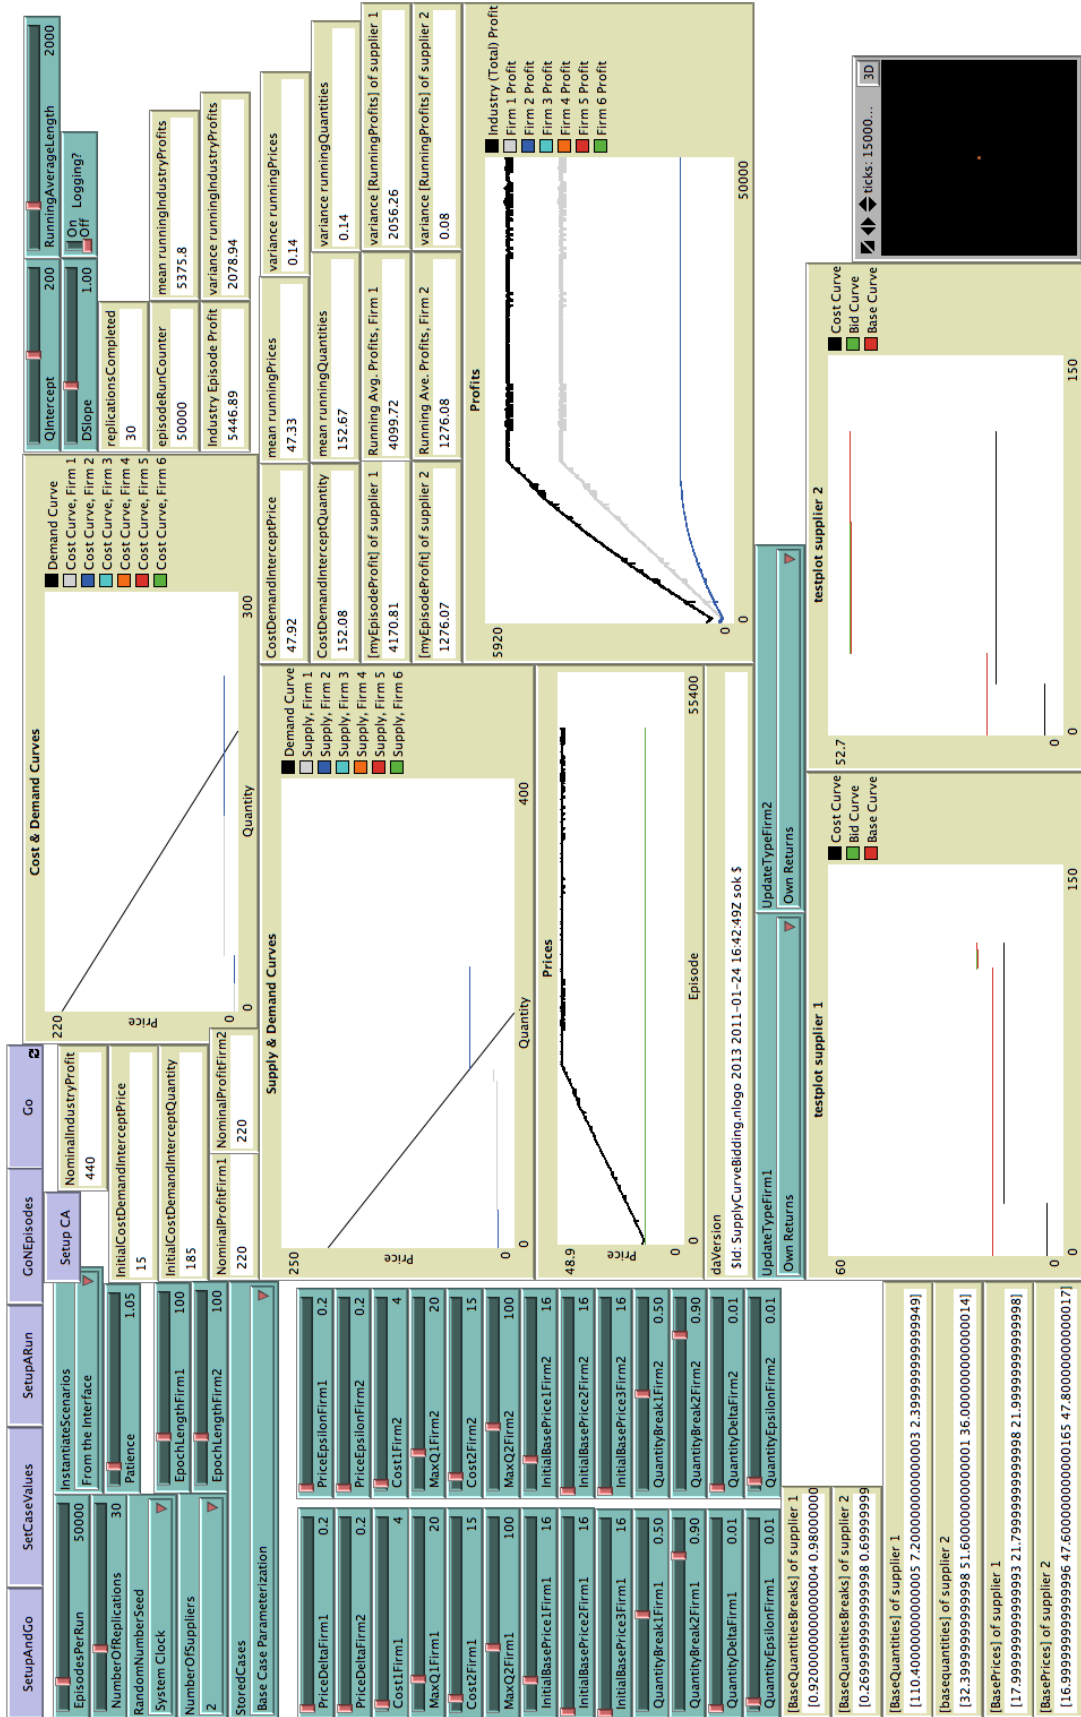


Figure 27: Base Case parameterization, but with low starting prices for all firms and segments.

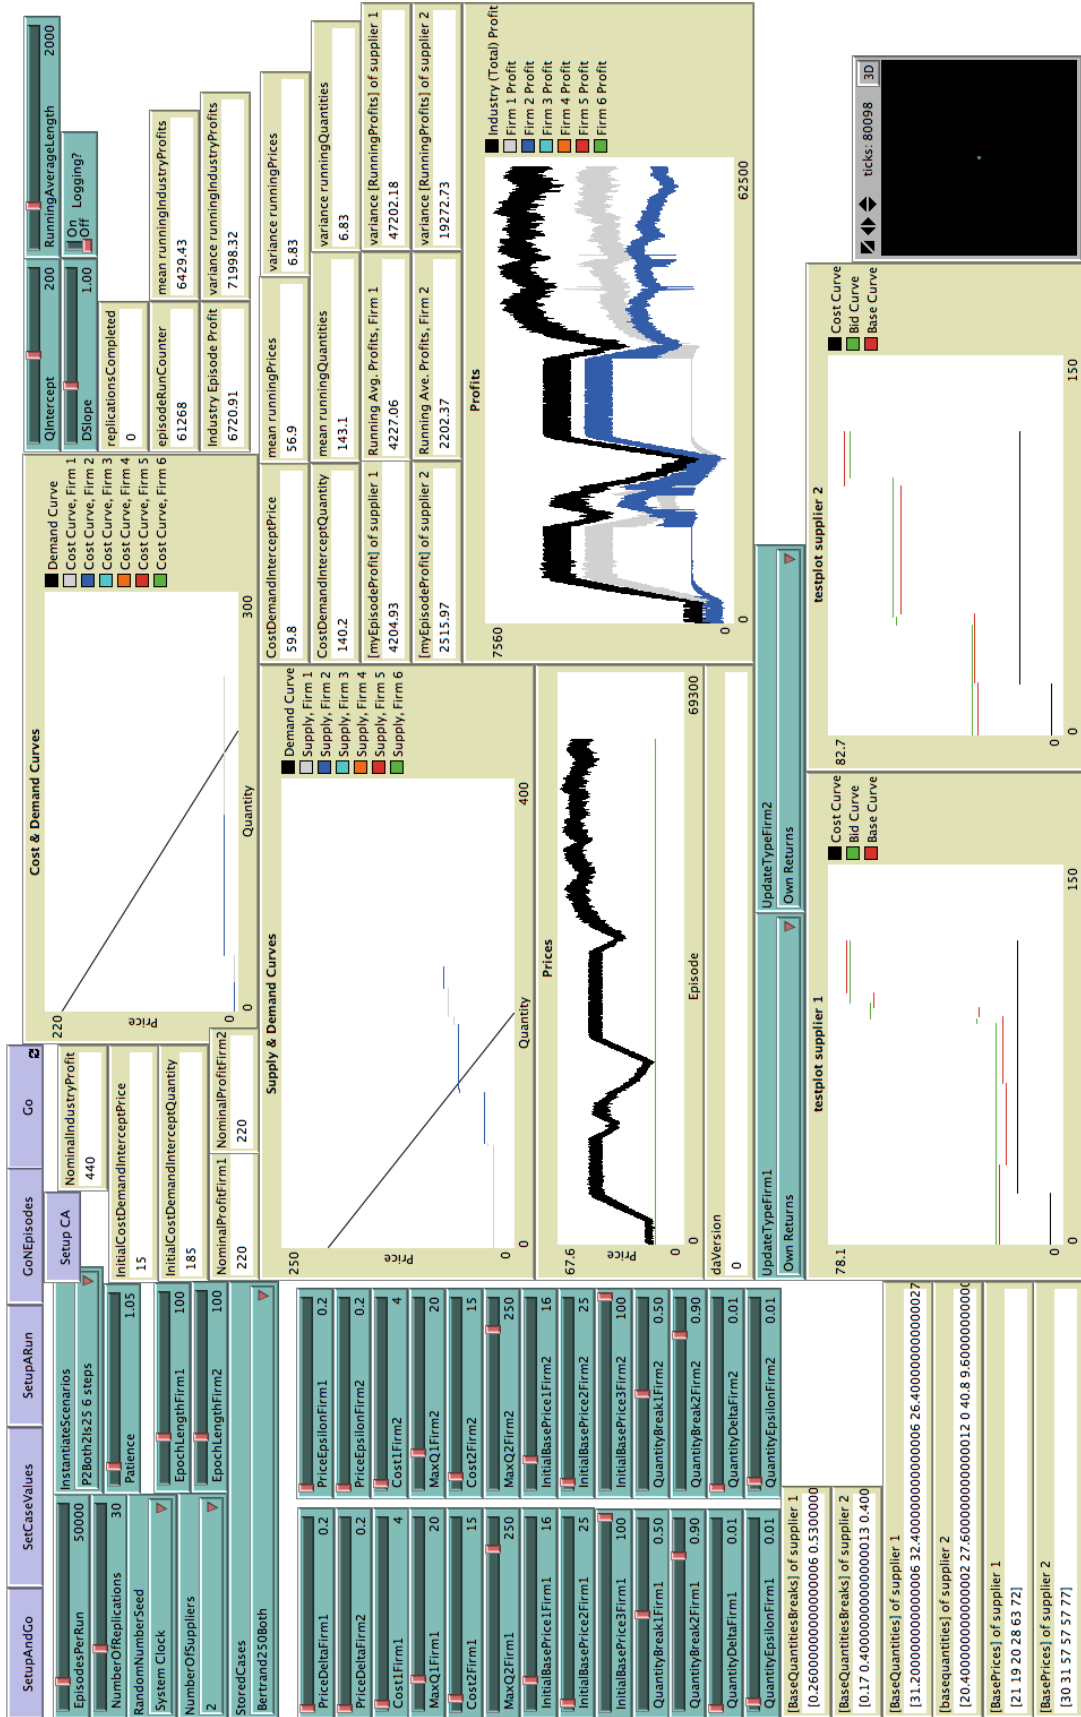


Figure 28: P2Both2Is25 6 steps stored case, which is P2BothIs25, but with 6 bid steps for each firm

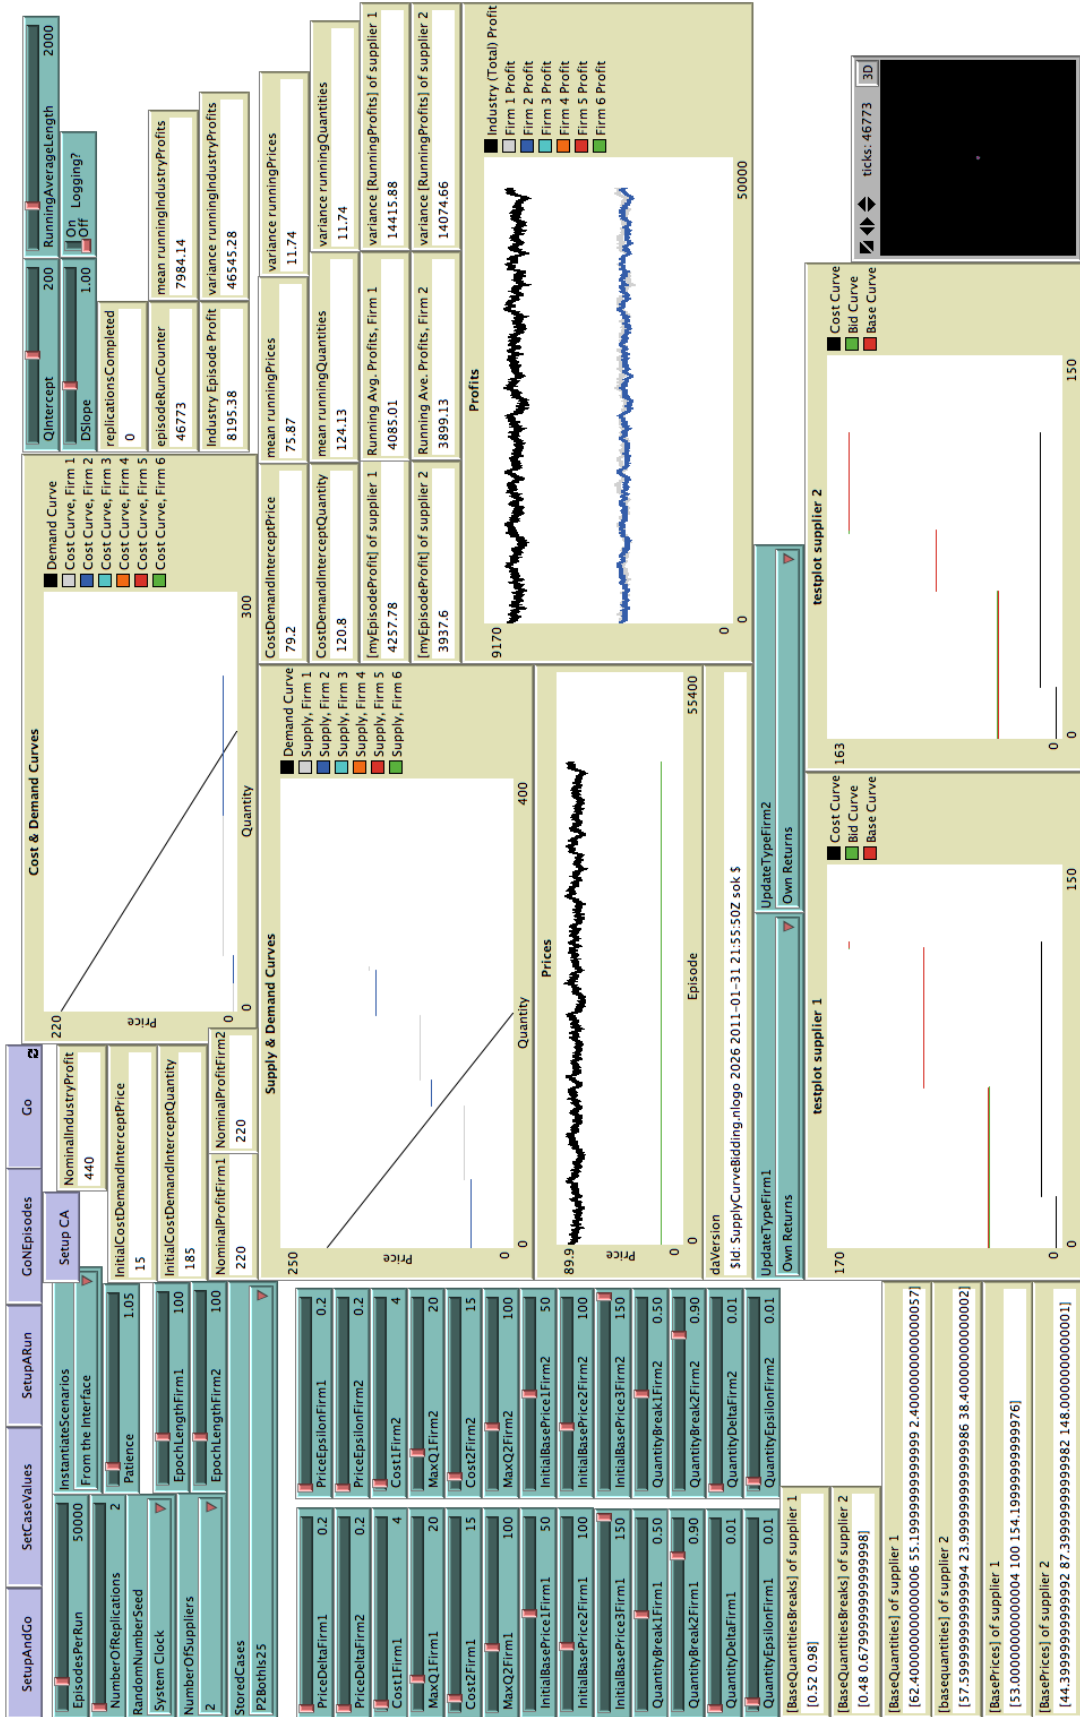


Figure 29: P2BothIs25 stored case, but both firms start bidding high

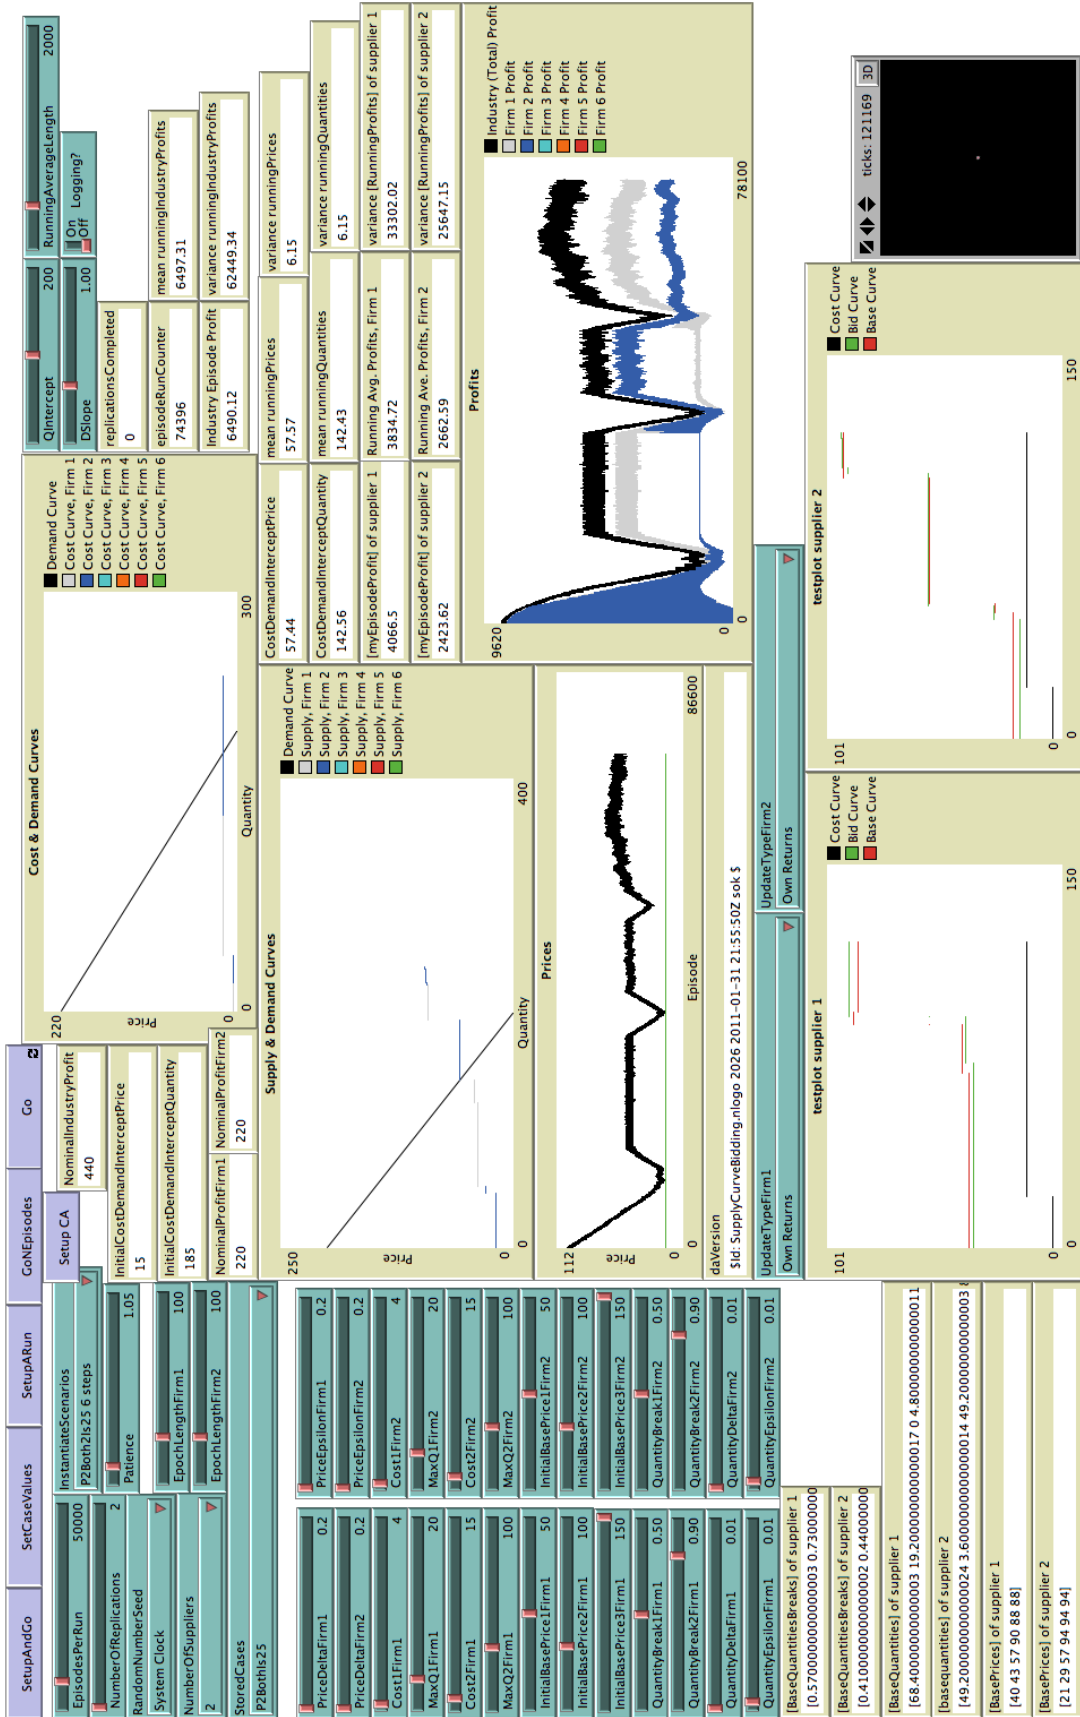


Figure 30: P2BothIs25 6 steps stored case, which is P2BothIs25 setup, but with 6 steps each firm, both firms start bidding high at 100 each step

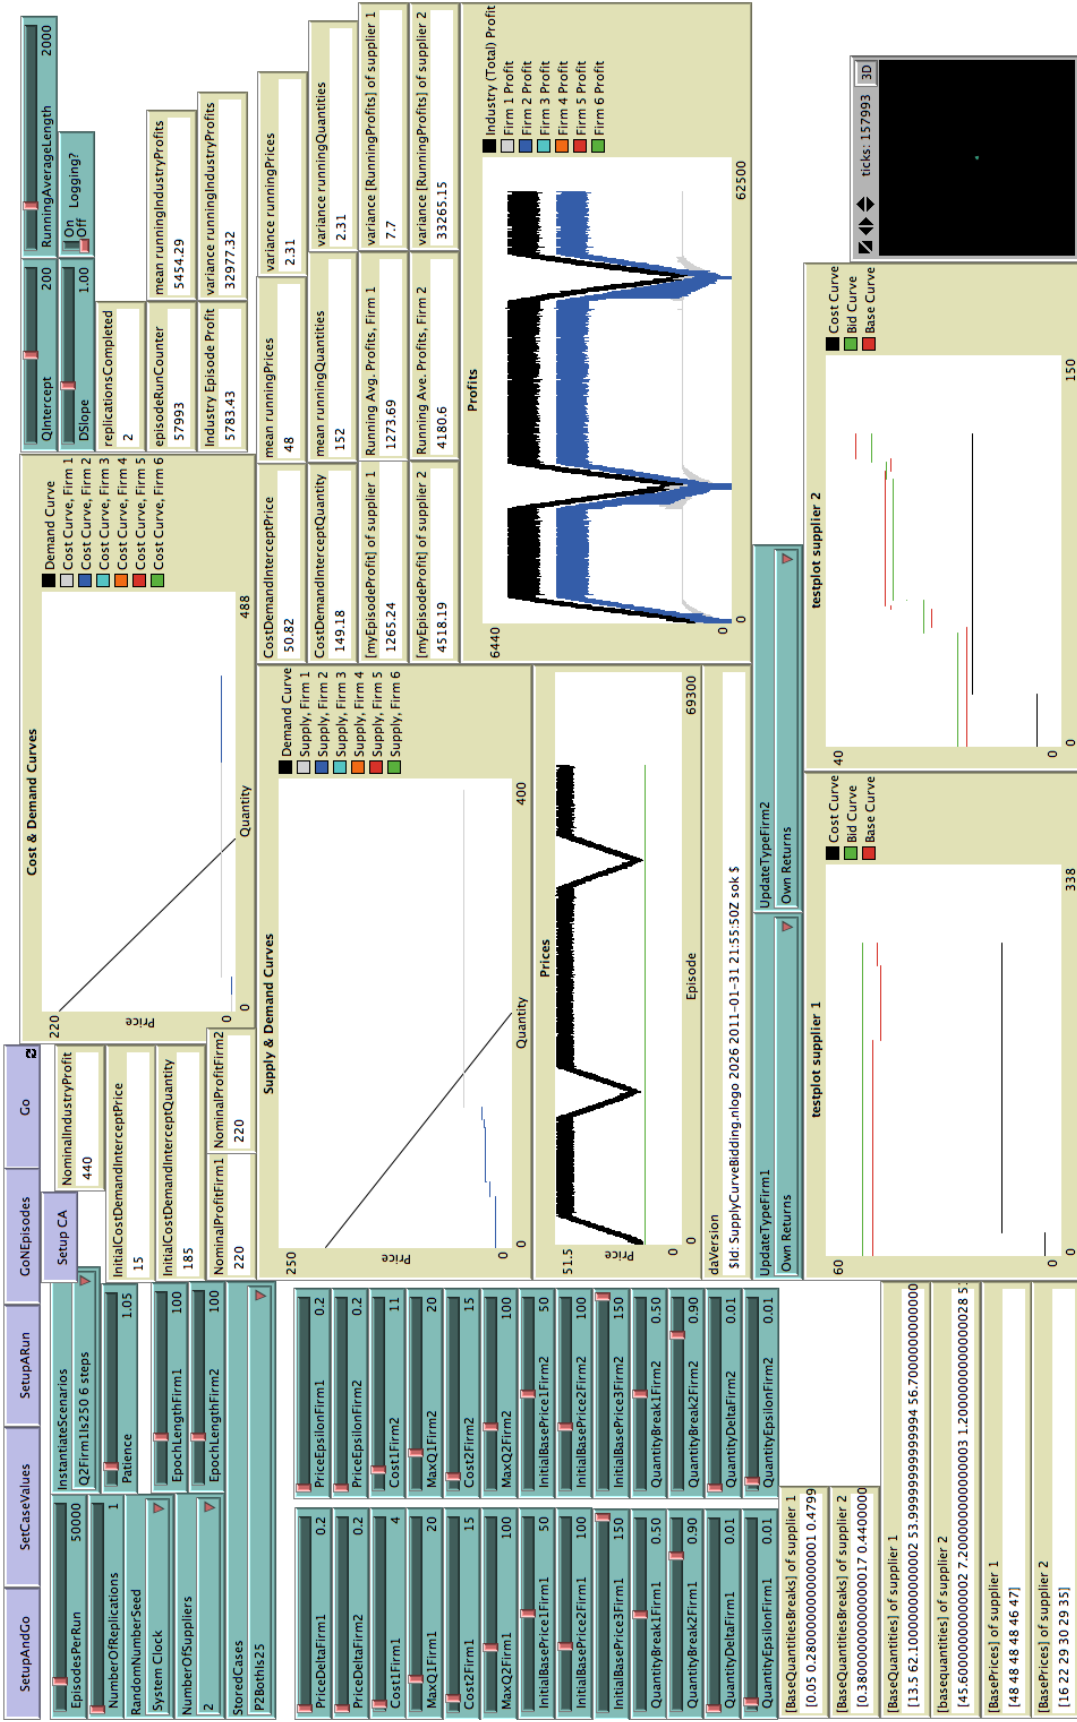


Figure 31: Q2Firm1Is250 6 steps stored case, which is the Q2Firm1Is250 setup, but with 6 steps for each firm, both firms start bidding low at 16 each step, a representative run

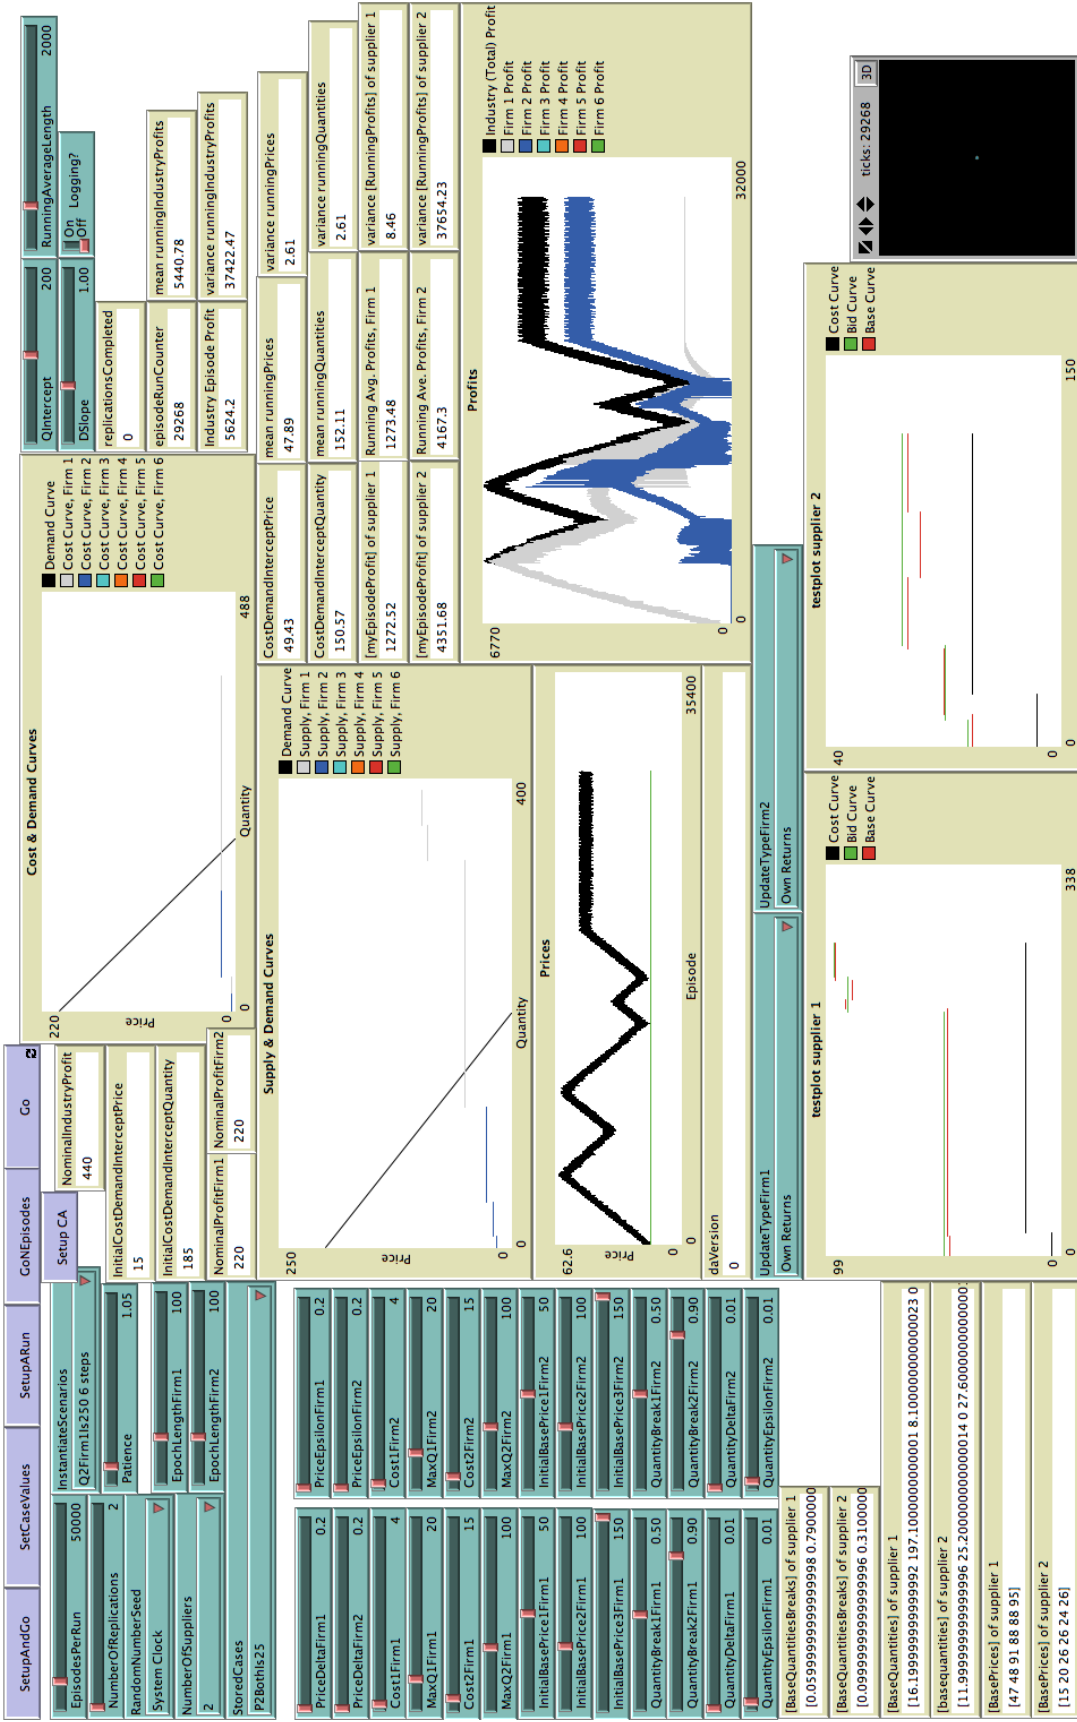


Figure 33: Q2Firm1Is250 setup, representative run (both firms update with Own Returns), but both firms have 6 bid steps. Firm 1 starts at base prices = 16, firm 2 at 100.

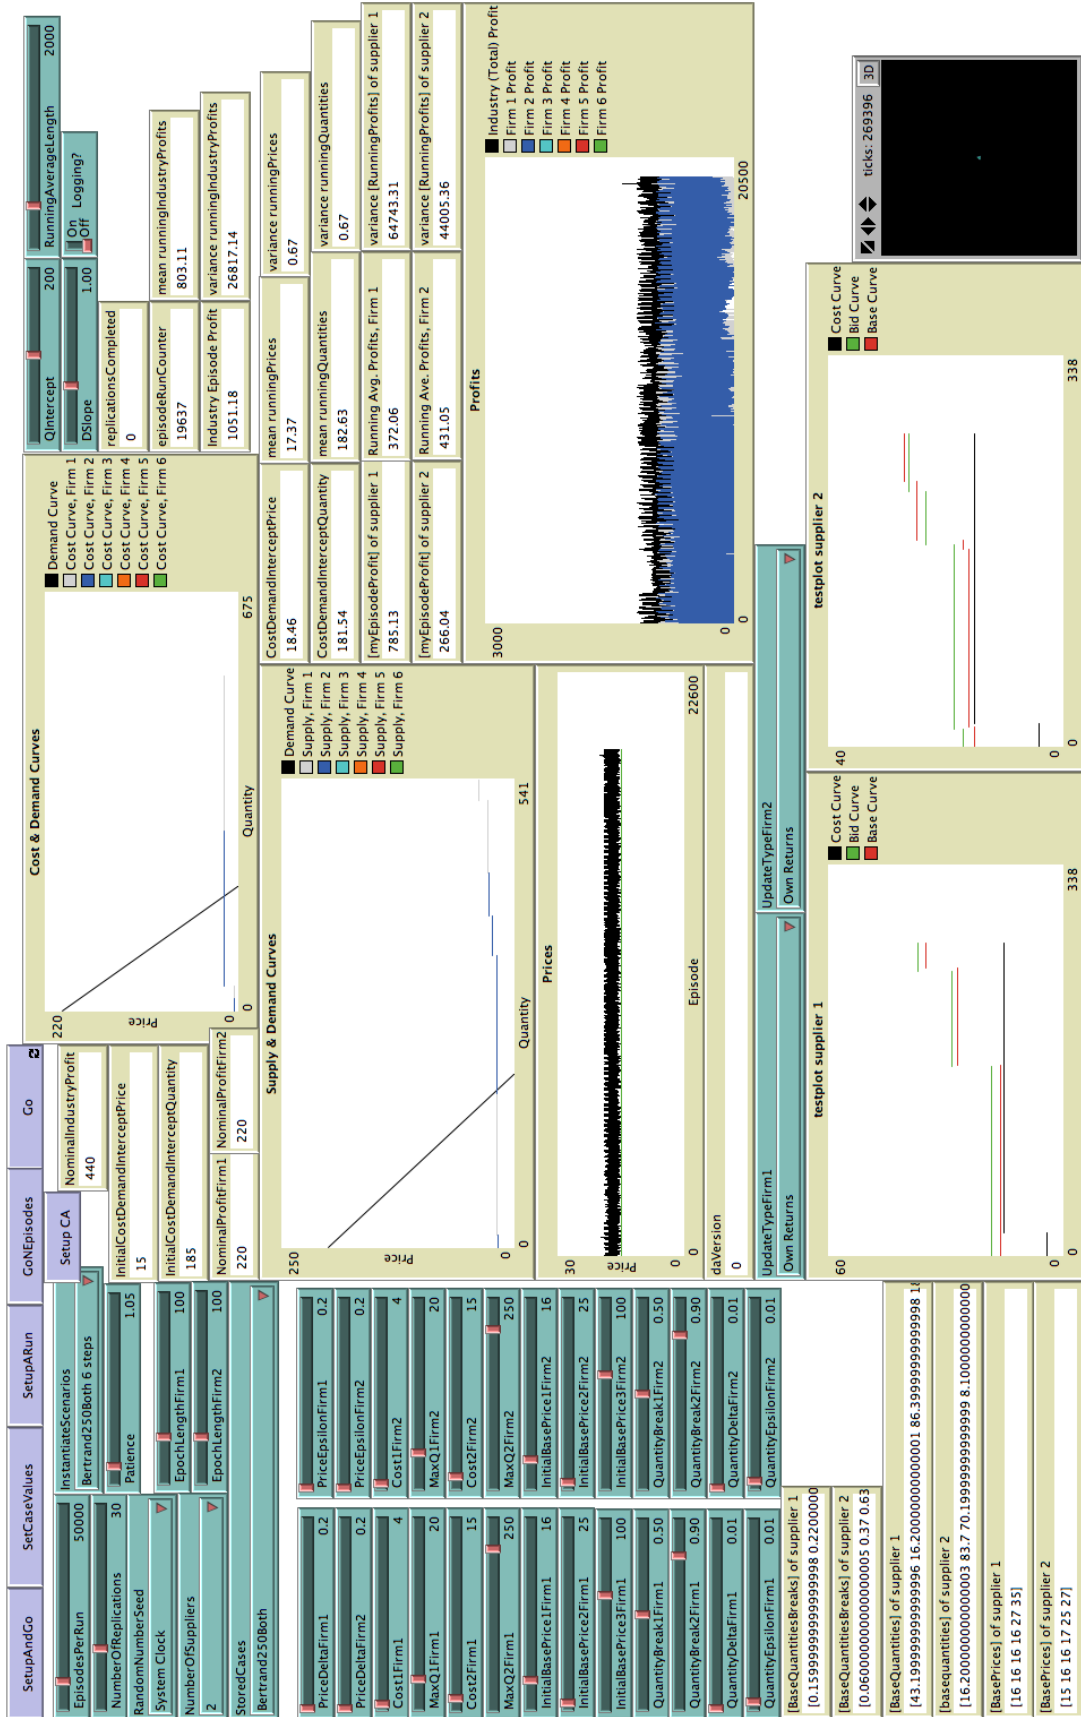


Figure 35: Bertrand250Both 6 Steps setup, representative run. Each firm has 6 bid steps. All firms start by bidding low.

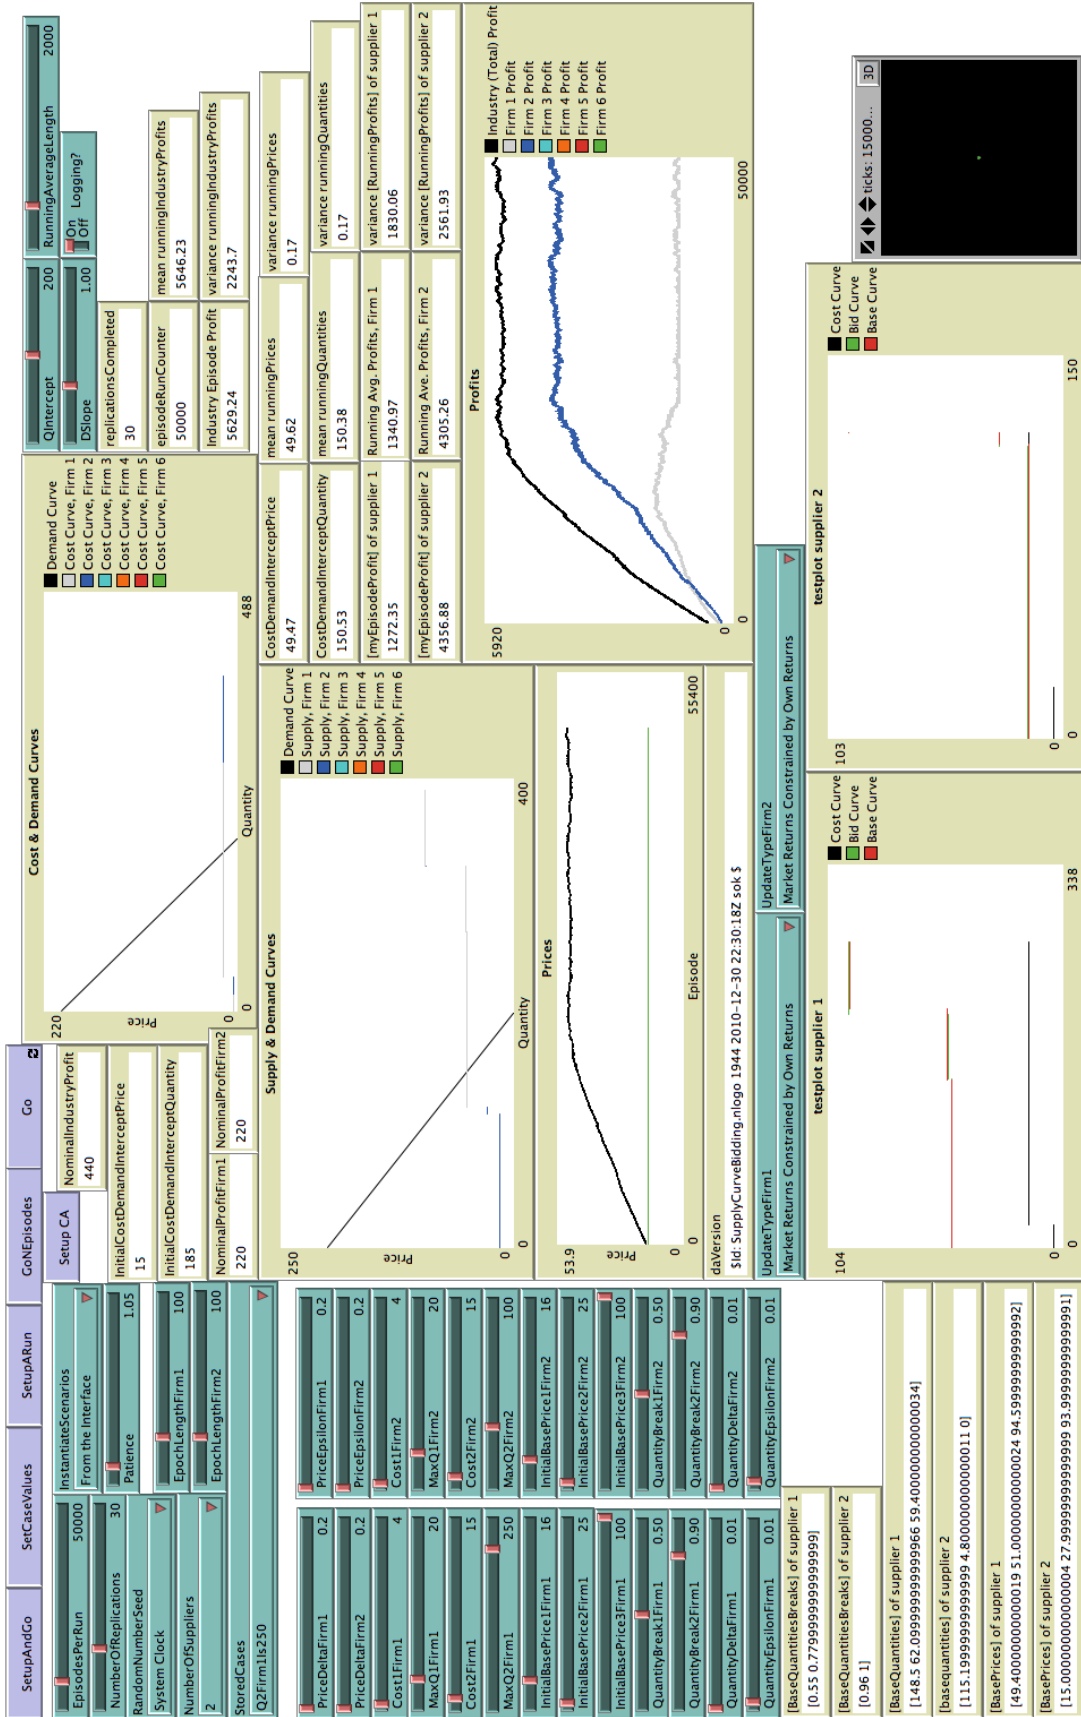


Figure 36: Q2Firm IIs250 Parameterization, representative run (both suppliers update with MR-COR)

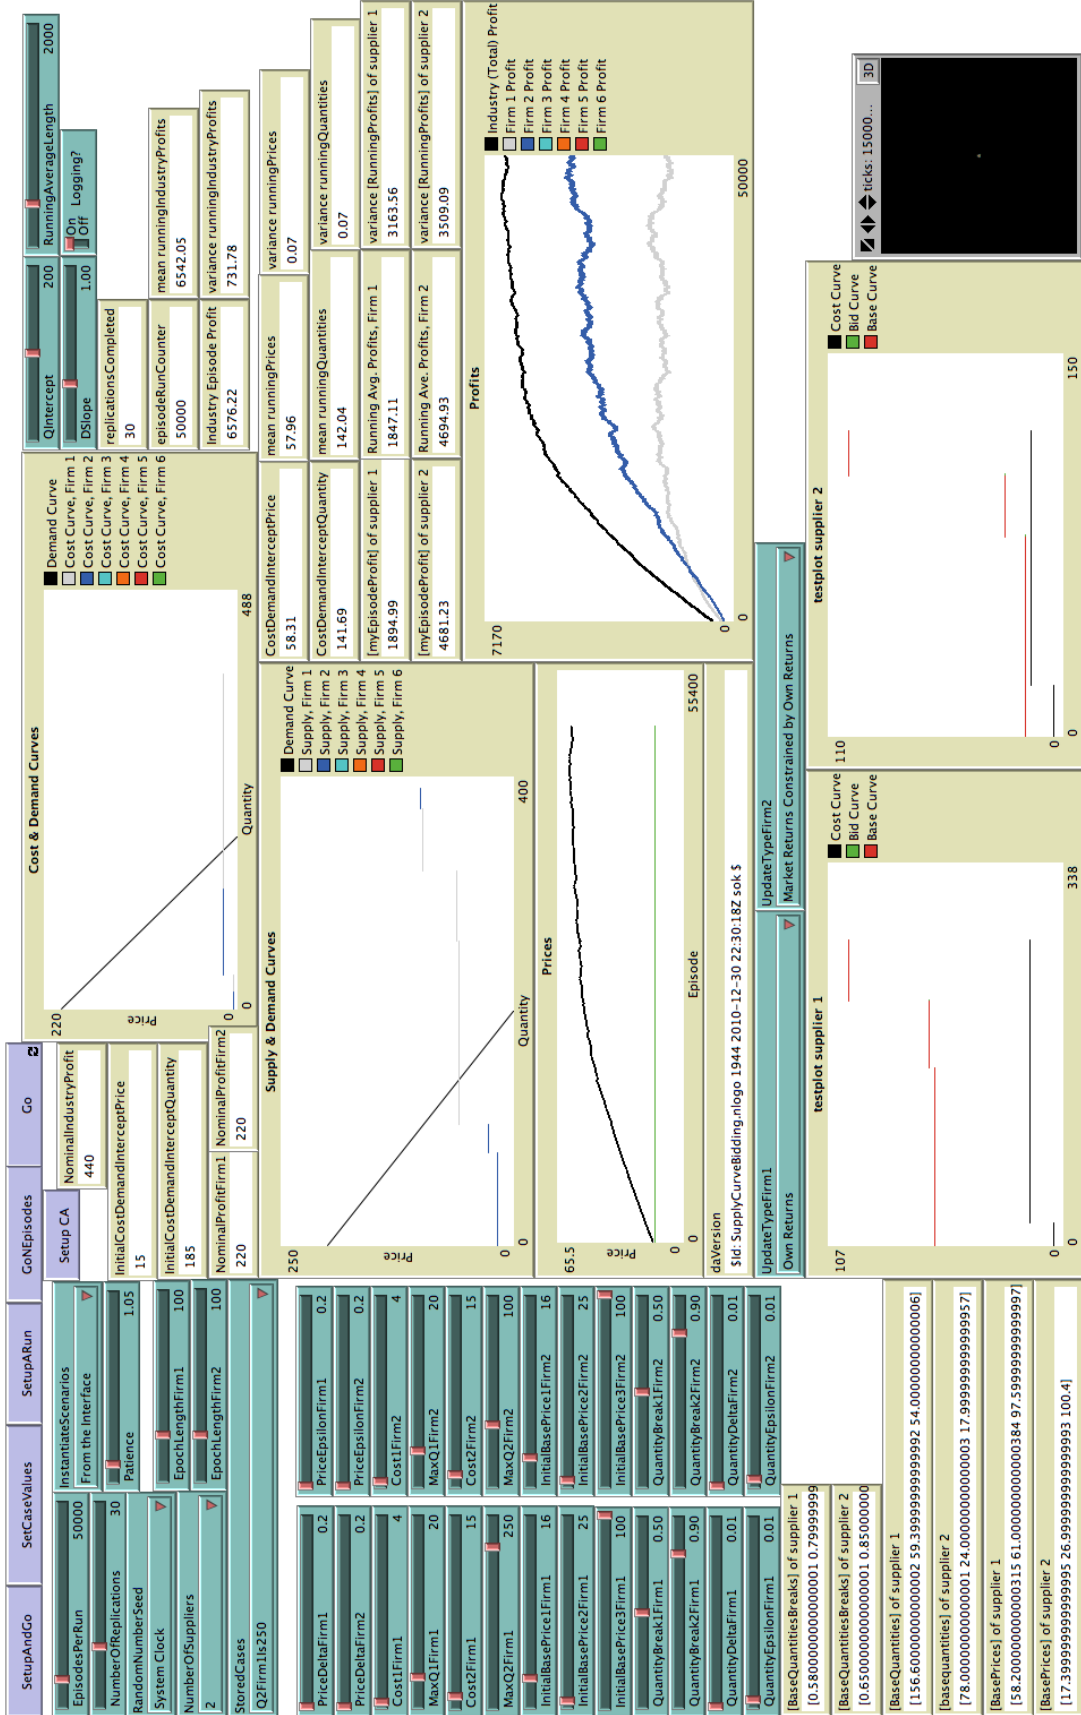


Figure 37: Q2Firm1Is250 Parameterization, representative run (S1 updates with Own Returns, S2 with MR-COR)

\$Id: Supplycurve_theorems_simulations_20110309supplementFred.tex 369 2011-03-10 20:21:51Z sok \$