Advance Production and Cournot Outcomes: An Experimental Investigation

Douglas D. Davis†*
Department of Economics
Virginia Commonwealth University
Richmond VA 23284-4000

January, 1999

Abstract

An experiment designed to assess the effects of advance production decisions on posted-offer market performance is reported. Six of the twelve triopolies were conducted under standard posted-offer rules. In the remaining markets sellers made binding production commitments prior to posting prices, an alteration that shifts the unique stage-game Nash equilibrium from the competitive to the Cournot outcome. As predicted, the advance production decisions raised prices and lowered output. Nevertheless, stable Cournot outcomes were never observed.

JEL Classification Codes: L1, C9

Keywords: Market Experiments, Cournot Competition, Advance Production

† Please send proofs, reprints and other correspondence to Douglas Davis, Department of Economics, Virginia Commonwealth University, Richmond VA 23284-4000.

* I thank without implicating Paul Brewer, Timothy Cason, David Harless, Charles Holt, Edward Millner Steve Perez, Steve Peterson, Charles Plott, Robert Reilly and Stanley Reynolds for helpful comments. Thanks are also due to Charles Holt for use of the laboratory facilities at the University of Virginia. Financial assistance from the National Science Foundation (grant SBR 9319842) and from the VCU Faculty-Grants-In-Aid Program is gratefully acknowledged. Data reported are available at FTP address fido.econlab.arizona.edu

1. Introduction

A long-standing debate in the industrial organization literature regards the use of Cournot or Bertrand specifications in modeling oligopolistic interactions. Bertrand-type price choices most closely match standard conceptions of inter-firm rivalry. However, given homogenous-products (and provided that capacity constraints are not binding) the competitive outcome is a unique Nash equilibrium for the Bertrand game, independent of the number of sellers. On the other hand, Cournot quantity-setting models, while institutionally artificial, generate more intuitively appealing price/cost markups that vary inversely with the number of sellers.

Kreps and Scheinkman (1983) potentially resolve this dilemma by showing that if sellers initially make binding production decisions, Cournot outcomes can be the unique Nash equilibrium for a pricing game.[1] Thus, Cournot outcomes may arise as the result of a long-run competitive process where firms compete on the basis of price after making capacity (plant-size) decisions. The Kreps and Scheinkman model occupies a prominent place in the industrial organization literature as a justification for using a Cournot specification. Thus, the behavioral consequences of the binding advance production commitments that drive results in this model represent an important open issue. Adherence to Nash predictions should not regarded as a foregone conclusion. Although Nash equilibria organize outcomes well in many contexts, Nash predictions fail persistently in several notable instances. For example, dominant strategy Nash equilibria fail persistently in both public goods environments and in noncooperative bargaining games (see e.g., Davis and Holt, 1993).

This paper reports an experiment that evaluates whether binding advance production commitments can move posted-offer market outcomes to Cournot predictions. Both Cournot predictions, and advance production decisions have been evaluated previously in the laboratory. An extensive experimental

[1] Davidson and Deneckere (1986) observe that the uniqueness of the Nash equilibrium prediction is sensitive to the buyer allocation rule used. See also Herk (1993).

literature evaluates Cournot predictions in matrix-game environments (e.g., Fouraker and Siegel, 1963 and Sherman, 1972, Feinberg and Husted, 1993; Holt, 1985; Mason, Phillips and Redington, 1991; Phillips and Mason, 1992). Mestelman, Welland and Welland (1987), and Mestelman and Welland (1988) examine the effects of binding advance-production decisions on posted-offer market performance. The novelty of the present investigation is that advance-production decisions are examined in a price-setting environment where (a) competitive, Cournot and collusive outcomes each have a reasonable chance of being observed *ex ante*, and (b) advance production decisions have unambiguous predicted effects. Combined, these features allow a test of the relevance of Cournot predictions in an institutionally plausible context.[2] Below, the experimental design and procedures are explained in section 2. Results are presented in section 3, followed by a discussion of results in section 4.

2. Experiment Design and Procedures

The effects of advance production are evaluated in a single-product triopoly with symmetric sellers. The identifiers printed below unit steps on the aggregate supply curve in Figure 1 illustrate cost conditions for the sellers, *S1*, *S2* and *S3*. As is evident from the figure, each seller may offer a maximum of six units, under conditions of increasing marginal costs. The first unit costs five cents, followed by a second unit costing 45 cents, third and fourth units costing 65 cents each, and fifth and sixth units costing 80 cents each. Sellers make production decisions in light of the step-wise linear demand curve labeled *D*, where the maximum willingness to pay for units drops in 20-cent increments from an initial high of 300 cents. The competitive price band ranges from 65 cents to 80 cents, and 12 units trade. The joint-profit-

[2] Binger, Hoffman, Libecap and Shachat (1990) also criticize the artificiality of the matrix-game quantity choice environment typically studied in the laboratory. To create a more institutionally-oriented decision structure, these authors had sellers make quantity decisions in light of an aggregate demand schedule. Wellford (1990) and Rassenti, Reynolds and Smith (1996) use a similar structure. Although more appealing than a matrix-game framework, this modified quantity-choice institution still deviates from natural market processes in that sellers never compete directly on the basis of price. Also unlike the experiment reported here, these papers do not evaluate performance relative to a baseline price-setting environment where the competitive outcome is a Nash equilibrium.

maximizing price is 200, with 6 units trading.

*** Figure 1 about here.***

A clean test of whether advance-production decisions can generate Cournot outcomes requires that advance production shifts the unique Nash equilibrium in a price setting game from competitive to Cournot outcomes. This issue is considered next. The section then concludes with a discussion of experimental procedures.

Equilibrium Considerations

The Posted-Offer Market: A version of standard posted-offer rules defines the trading institution for the baseline price-setting regime.[3] Markets consist of a series of trading periods. Each period starts by endowing agents with cost or value incentives. Trading then proceeds in a two-step sequence. First, sellers simultaneously post price and offer-quantity decisions. Penny increments define the price array, while single unit integer increments define the offer array. In this initial (no-advance-production) treatment sellers incur production costs only for units that sell. However, a seller must sell any unit offered, if a buyer wishes to purchase the unit.[4] After completing posting decisions a monitor publicly announces offer prices (but not quantities), and a shopping sequence commences in which a single simulated buyer makes all purchases possible without incurring a loss. The buyer purchases the lowest priced units available first. In the case of a tie the buyer divides purchases as equally as possible among the sellers posting the same price.

Given the cost and value incentives shown in Figure 1, the upper limit of the competitive price range, $P_e = 80$ uniquely defines the Nash equilibrium price for the posted-offer stage game. In the

---

[3] Posted-offer trading rules are among the most standard of laboratory trading-institution specifications, and a rich literature documenting posted-offer market results has evolved. Ketcham, Smith and Williams (1984) describe the institution carefully. Recent reviews of posted-offer experiments include Davis and Holt (1993), Holt (1995) and Plott (1989). The variant of posted-offer rules used here deviates from many implementations in that sellers possess as common knowledge full information about costs and demand.

[4] Sellers could offer units at a loss if they so desired. Although inconsistent with static profit maximization, multi-period trigger-strategies may include single period losses.

equilibrium per-seller earnings π_{ncp} = 140. Quantity allocations are not unique, and any aggregate supply in excess of 15 units defines an equilibrium allocation (implying that each seller offers at least of four units, and that no seller could sell more than a single unit by unilaterally deviating above the 80-cent price). To verify equilibrium existence, observe that under these conditions each seller earns 140 from the sale of their four infra-marginal units (= 4 x 80 - 5 - 45 - 65 - 65), plus zero from any marginal 80-cent units offered. Unilateral price cuts below 80 would reduce earnings, since only four units may be profitably offered at prices below 80. Unilateral price increases above 80 would, in the best case, result in the sale of a single unit, reducing revenues to 95.[5] Appendix A1 establishes uniqueness.

Advance-Production: Now modify posted-offer rules to include an advance production decision. The market proceeds as before, except that at the outset of each period each seller privately makes a production decision under the condition that all units offered incur production costs, regardless of whether or not the units subsequently sell. Following production decisions, a monitor publicly announces aggregate output and sellers make price and maximum offer-quantity decisions, as in the baseline posted-offer.[6]

To analyze the equilibrium for this game, start with the price-posting decision in the second stage. Given the first stage quantity decision, profit maximization equals revenue maximization, since production costs are sunk. In most instances the market-clearing price maximizes revenues.[7] Relevant

---

[5] Given residual demand of one unit, the profit maximizing deviator raises price to 100 cents (the next step on the demand curve), and sells a single low cost (five cent) unit. Residual demand of zero implies that any deviation reduces both sales and profits to zero. Notably, each Nash equilibrium at the competitive price is weak in the sense that some sellers must offer zero-profit fifth and sixth units. In the event no zero profit units are offered, no equilibrium in pure strategies exists for the stage game, and the equilibrium involves mixing. The multiplicity of cost and value steps impedes characterization of the equilibrium mixing distributions. Undoubtedly, however, most of the pricing density is close to P_e , since any price-triplet with more than one price above P_e, results in the exclusion of at least one seller from the market.

[6] Production and maximum offer quantity choices never differed. Although circumstances can arise where sellers elect to forego the sale of already-produced units (explained below), price increases accomplish these reductions more naturally than offer-quantity restrictions. Inclusion of an option to restrict offer quantities increased comparability with the posted-offer baseline.

[7] These allocations correspond to case (a) of Proposition 1 in Kreps and Scheinkman (1983), p. 335. The cases discussed

allocations include all quantity triplets where each seller produces four units or less, as well as all quantity triplets where aggregate output is ten or less. However, for sufficiently large aggregate output, residual demand at the market-clearing price is unitary elastic or inelastic, and the market-clearing price no longer uniquely maximizes total revenues. Residual demand at the market clearing price is unitary elastic for [aggregate, individual] quantity combinations of [11, 6], [12, 5], [13, 4] and [14, 3]. Residual demand becomes inelastic as individual output increases, holding aggregate output constant (e.g., [12, 6], [13, 5], [14, 4], etc.). For these allocations, the subgame equilibrium involves mixing over prices. Importantly, however, any seller anticipating an allocation that would make his or her residual demand inelastic at the market-clearing price would increase his or her earnings by cutting first-stage production (thus allowing him or her to earn the same revenue with lower production costs). In other words, no seller would ever voluntarily pick a quantity that would make it other than a dominant strategy to price at the aggregate output-clearing price.

Thus, equilibrium analysis for the game proceeds in terms of standard Cournot quantity-choices in the first stage. In the unique Cournot equilibrium each seller offers 3 units at a price $P_c = 140$. To verify that this allocation is a Nash equilibrium observe that for any seller a unilateral quantity increment from the Cournot allocation will result in the sale of an additional unit with a 65-cent cost. But the deviation increases aggregate quantity to ten and thus forces the market price down by 20, to 120. The incremental earnings loss of 60 (= 3[140 - 120]) from lower prices on the three units that would have sold at a higher price exceeds the incremental earnings increase of 55 (= 120 - 65) from selling an additional unit, making the deviation unprofitable. Larger quantity deviations are still more unprofitable, since they result in larger incremental revenue losses on units that would have previously sold (since prices fall yet further), and generate smaller incremental revenue gains from selling additional units (in addition to the price decrease recall that marginal costs rise to 80 cents for the fifth and sixth units).

---

below correspond to cases (b) and (c) of the same proposition.

Unilateral quantity decreases from the Cournot allocation also reduce profits. A unilateral decrease of one unit from the Cournot allocation raises the market price by 20 cents, thus increasing earnings on the two units the deviator still sells by 20 cents each, for a total increase of 40 cents. This incremental increase, however, is more than offset by the 75 cents lost by foregoing the sale of a third unit. (Recall, the third unit has a cost of 65, and previously sold for 140.) Further quantity reductions cause even larger net losses, since the deviator forgoes the sale of yet more profitable (lower cost) units in order to raise the price on the ever fewer units that are offered and sold. Thus, $P_c = 140$, $Q_c = 9$ is a Nash equilibrium.

Uniqueness of the Cournot allocation is relatively easily verified when costs and values are continuous. However, the non-convexities imposed by discrete cost and value incentives complicate the analysis, and can introduce a number of unintended asymmetric equilibria.[8] Alternative equilibria may be ruled out via consideration of payoffs for all three players in all possible quantity combinations, a cumbersome, but relatively straightforward task reported in appendix A2.

Procedures

The experiment consists of 12 sessions, with six sessions conducted in each treatment. To give sellers time to experiment with different price/quantity choices, all markets continued for at least 35 trading periods. After period 35 the market was randomly terminated by the result of a dice roll at the end of each period. The random termination rule was included to evaluate the effects of changes in the induced discount rate on the tendency of sellers to play trigger-strategy equilibria. However, no evidence of trigger-strategy play was observed in any market, and for brevity I omit discussion of this dimension of the experiment.[9] Experiment results in this regard parallel observations by Kruse (1993), who reports that static equilibrium predictions tend to organize outcomes far better than dynamic predictions in posted-

[8] This problem exists generally in discrete-choice environments. For example, Holt (1985) identifies several asymmetric equilibria that Fouraker and Siegel (1963) failed to notice in one of their discrete-choice games.

[9] An earlier version of this paper (available from the author) reports the effects of changing the termination rule.

offer market games.

All sessions were conducted under conditions of full and common-knowledge information about costs and demand. Each session commenced with an oral reading of instructions explaining the decision-making and record-keeping process. Instructions included information about the common cost structure and the aggregate market demand schedule (in tabular form). Following the instructions participants, seated in visually isolated booths made decisions, calculated profits and maintained their own earnings records.

The participant pool consisted of undergraduate students enrolled in principles and intermediate economics courses at the University of Virginia. All participants had each previously participated in a market experiment, though in a different environment and using different instructions. No one participated in more than one of the sessions reported here. Participant payments included a $6.00 appearance fee in addition to their salient earnings.[10] Additionally, participants started each session with an $8.00 (lab dollar) initial endowment, granted to offset any losses realized in early periods of the advance-production sessions. A 4 to 1 lab dollar/U.S. currency conversion rate kept the overall earnings level reasonable, while still giving sellers a reasonably thick grid of possible price choices. Participants knew the conversion rate as common knowledge. Earnings for the 1½ to 2 hour sessions ranged from $17.25 to $32.50, and averaged $23.50 (U.S.).

4. Results

The relatively minor procedural adjustment of submitting binding advance production decisions prior to posting price decisions affects market performance dramatically. Inspection of contract sequences for the first 10 trading periods of a representative posted-offer market (session PO3) and a representative

[10] The recruitment of alternates ensured enough participants appeared for each session. Unneeded alternates received a $10.00 flat fee.

advance-production market (session AP4) shown in Figure 2 offers some insight into the effect of the AP treatment on the contracting process. In the posted-offer market shown in the upper panel, price postings (hollow dots) decay in an orderly fashion toward the competitive prediction P_e. Contracts (crosses extending to right of the price posting) reflect the buyer's purchasing convention of buying units first from sellers posting the lowest prices. As is typical of posted-offer markets, posted prices decay across periods as the highest pricing seller in one period tries to edge under the postings of the other sellers in the subsequent periods.

*** Figure 2 about here ***

The advance production market shown at the bottom of Figure 2 exhibits much more volatile behavior. In the initial periods, seller confusion explains some of this volatility. Some sellers failed to appreciate that the sum of their individual production decisions implied the market-clearing price illustrated each period by the horizontal P_e line. In periods 1 to 4, seller prices (hollow dots) diverge uniformly from the (bolded line) market-clearing price. By period 6, however, the divergence between seller price-postings and the market clearing price falls substantially. Nevertheless, considerable volatility remains, as the sellers fail to repeat quantity decisions across periods.

The pricing patterns for the initial periods of the posted-offer and advance-production sessions in Figure 2 persisted throughout remaining periods of these sessions. The other posted-offer and advance production sessions also generated very similar outcomes, as can be seen by examining mean transactions price, quantity and average earnings series for the posted-offer ("PO") and advance-production ("AP") sessions in the left and right panels of Figure 3. In each panel the thin lines illustrate results for separate sessions, while the thick line summarizes the average outcome for the treatment as a whole.

*** Figure 3 about here ***

Comparing series across the panels in Figure 3 suggests the primary experiment results. First, the posted-offer series substantially convergence to competitive predictions. Price and earnings series, shown in the upper and lower panels, uniformly collapse on the respective competitive predictions of 80 cents and 140 cents. The quantity series, shown in the middle panel, reflect only slightly less complete competitive convergence, with average sales quantities in latter session periods hovering between the competitive prediction of 12 units, and the nearest feasible deviation of 11 units. Second, comparing across panels observe that the advance-production decision induces a substantial treatment effect. Perhaps most obvious are the increased volatility of individual sessions, and the increased dispersion of outcomes within treatments. But the figure also reflects some tendency for the markets to respond as predicted to the institutional change. Comparing the bolded overall-average lines across treatments, observe that the advance-production sessions tend to generate higher average prices, lower average quantities, and higher average earnings than the posted-offer sessions. Third, examining the advance-production series in light of Cournot predictions suggests no obvious tendency toward Nash predictions. Average prices and earnings tend to remain below Cournot predictions, while average quantities remain above the Cournot prediction.

Prior to a formal statement of results, the method of data analysis warrants some discussion. As is typical of market experiments, the data reported here follow a convergence process that is not well understood theoretically. Moreover, the dynamic process underlying many of the series likely includes serial correlation and heteroskedasticity. That said, a regression analysis that exploits all of the data generated has persuasive advantages. Most prominently, the analysis might cast some light on the convergence process in each treatment. In any case it avoids the problematic selection of some arbitrarily determined subset of the data (e.g., decisions from the last 10 periods) as reflecting “stable” decisions following whatever convergence process that may have occurred. In what follows we use an econometric specification adapted from Noussiar, Plott and Reizman (1995) to evaluate market convergence tendencies.

The specification evaluates the final average output for a treatment cell after letting each session start from unique initial values. Define t as the period number, and D_i as a series of dummy variables which take on a value of one for observations in session i, $i = \{1,2..6\}$ and zero otherwise. Then regression equation (1) estimates convergence behavior for output variable y:

$$y_{it} = \beta_1 D_1 \frac{1}{t} + \beta_2 D_2 \frac{1}{t} + \beta_3 D_3 \frac{1}{t} + D_4 \beta_4 \frac{1}{t} + \beta_5 D_5 \frac{1}{t} + \beta_6 D_6 \frac{1}{t} + \beta_{fin} \frac{t-1}{t} + u_{it} \qquad (1)$$

The weight $1/t$ places particular emphasis on the initial observations in a session, while the common weight $(t\text{-}1)/t$ places heavier value on the later observations. Thus, each "initial" β_i parameter estimates the starting value of an output series in session i while the "terminal" β_{fin} parameter estimates the terminal value of the output variable for all sessions in the treatment cell.

Data for the posted-offer and advance-production sessions are estimated separately. Estimated output variables, y_{it}, include mean transaction prices, mean sales quantities, and average earnings. The estimates allow for session-specific first order autocorrelation. White's (1980) method adjusts for heteroskedasticity.[11]

Regression results, summarized in Table 1, reflect salient features of the price, sales quantity and earnings series shown in Figure 3. For example, in the price series for the posted-offer markets, summarized in column (1) in the upper part of Table 1, each of the six initial parameter estimates exceed the essentially competitive terminal price estimate $\beta_{po\ fin}$ = 79.51 cents, as seen in Figure 3. Similarly, the sales quantity estimates in column (2) increase from heterogeneous but low initial values toward close to the competitive terminal value $\beta_{po\ fin}$ = 11.47, and the earnings estimates in column (3) generally reflect a decay to the essentially competitive terminal value $\beta_{po\ fin}$ = 141.01 cents. Note also that the $\bar{R}^2$ statistics for the posted-offer series range from .82 to .95, indicating that in all instances the regression equations for the posted-offer markets capture much of the movement in the data.

[11] Unlike Noussair et al. (1995) the data were adjusted by session-specific first order autocorrelation. This technique, suggested by Kementa (1987), pp. 618 -620, controls for the rather different adjustment patterns observed across sessions.

*** Table 1 about here ***

Regression results for the advance-production sessions, summarized in the bottom half of Table 1, similarly reflect salient features of the price, sales quantity and earnings series shown in right side of Figure 3. Most prominently, the $\overline{R}^2$ statistics for the advance-production sessions range from .13 to .60, indicating that the estimating equations capture considerably less of the movement in the volatile advance-production sessions than did the comparable equations for the posted-offer sessions. Initial and terminal parameter estimates for the advance-production series also convey information pertaining to movement within series. Comparison of initial to terminal price parameter estimates, in column (1), reflect a less pronounced decay of prices less than in the posted-offer markets. Moreover, the terminal price estimate, $\beta_{ap\ fin} = 95.44$ cents, while well above the competitive prediction, falls far below the 140-cent Cournot prediction. Comparison of initial to terminal estimates for the sales quantity series, summarized in column (2), reflects a tendency for sales quantities to increase throughout the course of advance-production sessions. However, the terminal sales quantity estimate $\beta_{ap\ fin} = 10.72$ although below the comparable posted-offer estimate, rises far above the Cournot prediction of nine units.

Comparison of initial to terminal earnings estimates in the advance-production sessions reflects a tendency for earnings to increase throughout the sessions, a dynamic opposite to that observed in the posted-offer sessions. The obligation of sellers in the advance-production sessions to pay for offered but unsold units, causes the difference in earnings patterns across institutions, because this feature increases the cost of initial errors. Nevertheless, the terminal earnings estimate for the advance-production series, $\beta_{ap\ fin} = 183.18$ cents, shown in column (3) resembles the terminal price and quantity estimates in that it deviates from the competitive outcome in the direction of the Cournot prediction, but remains far from the Cournot prediction of 300 cents.

The equation (1) regression specification very conveniently provides some objective criteria for evaluating the degree of market convergence to theoretical predictions. As suggested by Noussair et al.

(1995), a series can be said to exhibit *strong convergence* if the estimated β_{fin} does not differ significantly from the predicted value. Failing strong convergence, a *weak convergence* condition may apply. This condition presumes that initial choices for each series tend to be randomly distributed over the range of possible values, and thus will tend to deviate further from theoretical predictions than terminal values, when the theoretical predictions have some drawing power. A series satisfies weak convergence when a β_{fin} estimate is significantly closer to the theoretical prediction in absolute value than an initial β_i estimate. [12]

The superscripts appearing next to parameter estimates in Table 1 summarize results of the weak and strong convergence comparisons, evaluated at a 95% confidence level: a "*w*" aside a β_i estimate indicates that the series satisfies weak convergence to the relevant equilibrium prediction, while a "*" indicates that the weak convergence test does not apply, since initial and terminal estimates do not differ significantly. By default, no superscript by a β_i estimate indicates that the series fails weak convergence. Similarly, an "*s*" by a terminal β_{fin} estimate indicates that the treatment satisfies strong convergence, while no superscript indicates that the terminal estimate fails strong convergence.

Comparing β_{fin} estimates for each series establishes differences across treatments. The "*no*" subscript aside each of the β_{fin} estimates for the advance-production regressions indicate no overlap of 95% confidence intervals about the terminal value estimates for each advance-production series and its posted-offer counterpart. The results of these comparisons allows formal justification for each of the observations made at the outset of this section. Consider first the tendency for posted-offer markets to converge toward the competitive prediction. This calibration result represents a first finding:

---

[12] My weak convergence standard is more restrictive than that used by Noussair et al. (1995), who require only that β_{fin} be closer to the competitive prediction than β_i. The more restrictive standard used here introduces instances where weak convergence fails to apply. However, under this standard, sessions pass weak convergence only by rejecting the null hypothesis of no movement in the series.

Finding 1: *Posted-offer markets in this design converge toward competitive (Nash) equilibrium price, earnings and sales quantity predictions*.

Support: As shown in column (1) at the top of Table 1, the posted-offer price series passes weak convergence in each of the four instances where the test applies. Further, the treatment as a whole passes strong convergence. The posted offer earnings regression in column (3) generates similar weak and strong convergence test results, except that the weak convergence test applies in only two instances. The sales quantity series, shown in column (2) exhibits somewhat less complete convergence, since the terminal value estimate for the sales quantity series, at 11.47 fails strong convergence. Nevertheless, the competitive sales quantity prediction has considerable drawing power. The terminal quantity estimate deviates from the competitive sales quantity prediction by only a small amount. The upper bound of the 95% confidence interval about the terminal quantity estimate (11.61 units) misses the 12-unit competitive quantity prediction by only .39 units, or roughly 13% of the distance between Cournot and competitive predictions. Moreover, each of the six sales quantity series pass weak convergence.•

Consider next the effects of introducing the advance-production decision. Comparing terminal posted-offer estimates each for series with comparable advance-production estimates allows evaluation of this treatment effect. Results of this comparison, detailed below, form the second result.

Finding 2: *The addition of binding advance production decisions raises prices, lowers sales quantities and raises profits relative to competitive levels observed in baseline posted-offer markets*.

Support: Comparing β_{fin} estimates across treatments provides succinct support for this finding. As indicated by the "*no*" subscripts printed aside each of the β_{fin} at the bottom of Table 1, 95% confidence bands for terminal price, sales quantity or earnings estimates do not overlap across treatments. The advance-production treatment yields significantly higher terminal prices and earnings, and significantly lower terminal earnings than the posted-offer treatment.•

Despite the movement of prices, sales quantities and earnings in the advance-production sessions away from competitive outcomes and in the direction of Cournot predictions, the advance-production markets do not come close to generating Nash predictions. This is the third finding.

Finding 3: *Advance-production markets in this design do not converge toward the Cournot (Nash) equilibrium prediction. Prices and earnings are significantly below Nash predictions, and sales quantities are significantly above Nash predictions.*

Support: The price, quantity and earnings series estimates for the advance-production sessions, shown in the bottom half of Table 1, provide support for finding 3. Start with the price series. The 95.44 cent β_{fin} estimate for the price series, summarized in column (1) falls far below the 140 cent Cournot prediction. The absence of an "*s*" superscript next to the estimate indicates that this difference is significant at a 95% confidence level, and that the series thus fails strong convergence. Further, the terminal value estimate deviates from the Cournot prediction by a large margin. The upper bound of a 95% confidence interval about the terminal estimate (99.38 cents) misses the 140-cent Cournot prediction by more than 40 cents, or 67% of the range between competitive and Cournot price predictions. Note additionally that the weak convergence test fails in each of the four cases where the test applies.

Cournot predictions similarly fail for the advance-production sales-quantity and earnings series, and in each instance by large margins. For the quantity series, summarized in column (2) of Table 1, the lower bound of a 95% confidence interval about the 10.72 terminal sales-quantity estimate (10.51 units) misses the Cournot prediction by 1.51 units, or one-half the distance between Cournot and competitive predictions. For the earnings series summarized in column (3), the upper bound of a 95% confidence interval about the 183.18 cent terminal earnings estimate (196.23) misses the Cournot prediction by 103.77 cents, roughly two-thirds of the distance between Cournot and competitive predictions.

For the sales-quantity series, weak convergence test results provide further evidence of the failure of Cournot predictions, with weak convergence failing in five of the six instances. Notice, however, that the earnings series for the advance production treatment pass the weak convergence standard in all five instances where the test applies. In this case, the weak convergence standard is a bit misleading. Rather than indicating some tendency for convergence to Nash predictions, the advance production earnings series pass weak convergence only because sellers tend to recover from the very low (and occasionally negative) earnings arising from costly initial errors.•

4. Discussion

Advance-production clearly affects performance. The addition of a binding advance-production decision moves price, earnings, and quantity outcomes toward Cournot predictions. Nevertheless, advance-production markets adjust imperfectly. In contrast to the uniform stability of decisions about the competitive prediction in posted-offer triopolies, advance-production markets generate outcomes that both persistently deviate from Nash predictions and remain highly variable.

What explains the differential performance of unique Nash predictions in the alternative environments? Each of the alternative equilibria are unique in their respective environments, and no obvious refinement distinguishes them. One candidate explanation lies in the comparative complexity of the advance-production game. Perhaps behavioral convergence to equilibrium predictions in this context requires more than 35-40 trading periods. As suggested previously by the sequence of contracts for the first 10 periods of session AP4, shown in Figure 2, participants at the outset of sessions do not appreciate the generally mechanical nature of the pricing decision, given the quantity choice. As shown in row (1) of Table 2, sellers as a group deviate pervasively from the optimal price in periods 1-5: 78.8% of price choices were either above the market-clearing level (resulting in produced but unsold units) or were sizably (more than two cents) below the optimal price. Pricing decisions, however, did improve as the sessions progressed. Continuing across row (1), notice that by periods 20-25, "sizable" deviations represented only 20.8% of price choices, and that for the last 10 periods common to all sessions, the incidence of deviations fell to roughly 15.4%.[13] Despite the improvement in pricing decisions, the poor initial decisions may have made the early periods relatively uninformative. Possibly another 20-30 trading periods would generate stable outcomes.

However, the similarity of the results reported here to results of other experiments conducted explicitly in quantity-setting environments suggests that added complexity is not the primary explanation for the continued instability. Persistent variability is a prominent characteristic of data in many Cournot experiments. (See, for example, Fouraker and Siegel, 1963; Wellford, 1990; Phillips and Mason, 1992; and Binger, Hoffman, Libecap and Shachat, 1990). Perhaps most pertinent to the present study is Rassenti, Reynolds and Smith (1996), where despite relatively thick (5-seller) markets, and comparatively long sessions (60 periods each), output decisions remained both between Cournot and competitive predictions on average, and persistently volatile. Thus, I conjecture that the instability of the advance production environment is not due to the failure of the institutional framework to induce Cournot incentives, but rather to the instability of Cournot incentives themselves

An understanding of behavior in experimental Cournot markets remains elusive, and the topic remains an important area for continued research. Any attempt to thoroughly assess such behavior here would lead us rather far astray from our current focus.[14] Nevertheless, the following three observations provide may prove useful.

---

[13] The choice of two cents as defining a "sizable" price deviation reflects a balancing of sorts. On the one hand, any negative price deviation results in some foregone earnings, and may therefore be considered irrational. On the other hand, the costs of small negative price deviations are small. Such deviations never result in unsold units, and they do not affect best responses. (To affect best-responses unilateral negative price deviations, must exceed four cents.) For completeness note that use of a negative four-cent cut-off generates deviation rates virtually identical to those shown in Table 2. Use of a zero-cent cutoff generates results to those illustrated in the initial 5-period sequences, but deviation rates decay only to about 30% in periods 26-30 and 31-35. Finally, note that the calculations in column 1 of Table 2 exclude periods where posting the market clearing price was not a dominant strategy (e.g., equilibrium pricing behavior involved mixing).

[14] For further analysis, see Rassenti, Reynolds and Smith (1996).

First, participant response patterns do not conform to simple standard theoretical specifications. Perhaps the simplest adjustment dynamic is "best-response" play, where each seller makes output decisions that optimize earnings relative to their rivals' current choices. Cournot (1960) introduces this simple adjustment process to demonstrate the dynamic stability of quantity-choosing duopolists. However, in Cournot markets with more than two players, best responses are not always dynamically stable. As shown by Theocharis (1960) and Fisher (1961) unilateral best-responses to a non-optimally small aggregate output can cause sellers to collectively overshoot the equilibrium and vice-versa, causing cycling, or even explosively divergent swings in aggregate output.[15] Other response specifications damp the collective tendency of sellers to overshoot the mark and can be dynamically stable even when best responses fail. Alternatives include a "partial adjustment" dynamic, where output decisions are a weighted combination of the current play and the best-response to others' output decisions; and "fictitious play," where output decisions are based on the frequency of plays by others.

Observed output decisions doubtfully conform to any of these alternatives. Best-response play can be ruled out, since participants neither make best-response decisions initially, nor learn to do so in later periods. As shown in row (2) of Table 2, 95.5% of the decisions made in periods 1-5 deviated from best-responses. By periods 30-35, the percentage decayed only marginally to 82.1%. An array of partial adjustment and fictitious play specifications are possible, and evaluating directly all these alternatives is infeasible. However, persistent deviations from best responses in a given direction (either positive or negative) are inconsistent with dynamically stable variants of either alternative. As seen in the listing of positive deviations from best-responses, shown in parenthesis below the best-response deviations, in row (2) of Table 2, deviations are almost uniformly positive. In the first five periods, 94.4% of total decisions

---

[15] In the design examined here the stability of a best-response dynamic depends on the sellers' initial quantity choices. As may be verified by evaluating best-responses to each of the (6x6x6=) 216 possible initial choices, convergence to the Nash equilibrium occurs in 55 instances. Best responses to the remaining initial choices generate cycling behavior, with symmetric cycles between {2,2,2} and {4,4,4} in 116 instances, and asymmetric cycles between {2,3,3} and {3,4,4} occurring in the

deviated from the best-response in a positive direction. Although the incidence of positive deviations diminishes somewhat over time, it never falls below 65.6%, and for periods 30-35, is 70.3%. Thus, rather than making some sort of optimizing adjustment relative to past play, players appear to consistently produce too much.

*** Table 2 about here ***

A second observation regarding dynamic performance is that strategic considerations may explain the observed overproduction. Although participants do not exploit the repeated structure of the markets to enhance earnings via trigger strategies, some players do appear to repeatedly make aggressive (large) output decisions in a strategic effort to increase relative earnings. Efforts by a seller to "bully" the others into accommodating disproportionately high quantity choices can increase not only average output, but the volatility of decisions as well. Quantity peaks arise as the other sellers resist overproduction by the bully. Quantity troughs occur when the bully, and everyone else yield simultaneously.

But strategic efforts of this type often succeed. Seller earnings as a percentage of total earnings in his or her market are highly correlated with the seller's output as a percentage of market output. For periods 26-35, Spearman's rank order correlation coefficient for the relationship between relative earnings and relative output, $\rho = .745$. In fact, some participants managed to increase earnings not only relative to others, but relative to Nash predictions. Two of the three sellers who enjoyed above Nash earnings in periods 26-35 produced the most relative to their markets. On the other hand, such strategies do not always succeed. If competitors match aggressive quantity plays, earnings for all sellers fall. Perhaps not surprisingly, the markets with the highest average output generated the lowest average earnings. In periods 26-35 market earnings and market output correlate perfectly (Spearman's $\rho = -1$). Rassenti, Reynolds and Smith (1996) observe similar behavioral patterns.[16]

---

remaining 45 instances

[16] The frequent failure of such strategies is unsurprising. In a game of indefinite length, rivals would most likely emulate

Third (and not inconsistent with the two observations made above) Nash behavior may simply be too much to ask for in this environment. Other solution concepts may be more reasonable. Consider, in particular, the notion of *rationalizability*, developed independently by Bernhiem (1984) and Pearce (1984). Actions are said to be rationalizable if players make a best response to *some* rational belief about the actions of the others (rather than the Nash assumption that players make a best-response to a belief that the other players will make an equilibrium response). The relevant strategy set may be isolated by finding those strategies that survive strictly-iterated dominance.[17] In many instances the set of rationalizable outcomes and Nash equilibria overlap. In particular, in the posted-offer game, the competitive outcome is both the unique Nash equilibrium and the only rationalizable strategy. However, in the Cournot game, strategies {2,3,4} survive iterated strict dominance.[18] As is clear from the incidence of deviations from rationalizable outcomes, listed in row (3) of Table 2, participant decisions increasingly fall within the set of rationalizable strategies as the sessions progress. Although 38.9% of individual output decisions fell outside the set of rationalizable outcomes in periods 1-5, the percentage of deviations fell steadily, to 6.7% and 4.4% in periods 26-30 and periods 30-35, respectively. Assessing the importance of rationalizability as a means of organizing market behavior merits further investigation, as, more generally, do efforts to explain behavior in Cournot environments.

In closing, I mention a related issue for future research, that calls on the role of experimentation as a means of institutional design in addition to its role in theory evaluation. Although the advance-production institution does not generate stable Cournot outcomes, stability might be induced in related variant institutions. For example, stability may be enhanced with the addition of a second quantity-setting stage prior to price-setting where sellers could renege on part or all of their first-stage quantity choices.

---

profitable choices, rather than be deterred by them. Vega-Redondo (1996) shows that Cournot outcomes are evolutionarily stable when rivals copy the choices of the seller earning the highest profits

[17] With three or more players, the set of strategies that survive strictly iterated dominance satisfy *correlated rationalizability* Pearce (1984)

Plott, Sugiyama and Elbaz (1995) study such an institution in a natural monopoly context. Investigation of performance in these related institutions would not only help identify where stable Cournot outcomes might arise in naturally occurring contexts, but might also help improve the performance of markets that effectively operate on the basis of quantity choices.

[18] Review of payoff table A1 allows verification of this result.

Appendices

A1. Uniqueness of the Nash Equilibria in the Posted-Offer Treatment

Ruling out all price triplets other than [80, 80, 80] establishes uniqueness in prices and earnings for the Posted-offer treatment. Let us first eliminate triplets consisting of alternative common prices. At any common price below 80 each seller sells the same four units that would sell at 80, and thus each seller has an incentive to increase price unilaterally. At any common price p between 81 and 100, each seller sells 3.67 units in expectation (one-third of the aggregate 11 units that will sell). A unilateral price reduction of a penny to $p - 1$, implies that the fourth unit, costing 65 cents, sells with certainty, along with the fifth and sixth units, costing 80 cents each. Posting price $p - 1$ increases profits as long as the expected sales gain exceeds the loss on expected sales at the common price, or as long as

$$.33[(p - 1) - 65] + 2[(p - 1) - 80] > 3.67(1). \qquad \text{(a1)}$$

Solving, $p > 80.4$, which implies that any common price in the 81 - 100 cent range invites profitable deviation. Next, rule out common prices of 101 and above by noting that the reduced aggregate sales associated with higher prices simultaneously increases the left side of (a1), and decreases the right side of the equation, making deviation still more profitable. The left side increases for two reasons: First, the price-cost spread (the bracketed terms) increases as prices rise. Second, as prices rise expected per-seller sales at the common price fall, implying that the gain from a unilateral deviation increases. The right side of (a1) decreases with price increases, due to the fall in expected sales.

Now rule out price triplets consisting of asymmetric prices. Any triplet containing a price below 80 invites profitable deviation to 80 by the seller(s) posting the low price(s). Any triplet containing one or more prices above 80 implies that the high-pricing seller will be left out of the market. This seller increases earnings by pricing below one or both of the other sellers. Finally, uniqueness in prices rules out equilibria involving mixing over prices.

A.2 Uniqueness of the Nash Equilibrium in the Advance-Production Treatment.

Verify uniqueness of the Nash equilibrium in the advance production treatment by considering the profitability of deviations for seller, from every possible quantity outcome. The eighteen 6x6 panels in Tables 4, 5 and 6 summarize relevant earnings information. The panels are divided into six three-panel rows. Each table contains two rows. Moving from left to right along each row, the panels report earnings for *S1*, *S2* and *S3*, respectively. Moving down the tables, each row of panels reports earnings under a different output choice by *S3*. Finally, within panels, the 36 entries summarize earnings possibilities for a particular seller-type, for each of the six possible quantity choices available to *S1* (columns) and *S2* (rows). Thus, for example, the left-most panel in the top row of Table 3 lists earnings possibilities for seller *S1*, for all possible quantity choices for *S1* and *S2*, given that *S3* produces one unit. Similarly, the right-most panel in the bottom row of Table 5 lists earnings possibilities for seller *S3*, for all possible quantity choices by *S1* and *S2*, given that *S3* produces six units.

To assist in the identification of equilibria, best-responses for *S1*, *S2* and *S3* are highlighted in the panels. In the left-most column of panels, vertical lines to the left of entries highlight the best responses of *S1* to choices by *S2* and *S3*. Multiple vertical lines in a row indicate indifference over quantities for *S1*. Similarly, in the middle column of panels, dashed lines highlight the best responses of *S2* to choices of *S1* and *S3*. Multiple dashed lines in a column indicate indifference over quantities for *S2*. In the right-most column of panels, shaded areas highlight best responses for *S3*. Seller *S3* responds to choices by *S1* and *S2* by moving up or down the row panels. Areas shaded in more than one of the right-most panels indicate indifference over quantities for *S3*. For reference, the right-most column of panels also repeats the highlights for *S1* best responses (vertical lines) and *S2* best responses (dashed lines).

Intersections of best-response highlights for all seller identify Nash equilibria: As seen in the tables, only the allocation [3, 3, 3] contains shading, a vertical line to the left and underlining.

References

Bernheim, Douglas, 1984, Rationalizable strategic behavior, Econometrica 52, 1007-1028.

Binger, Brian R., Elizabeth Hoffman, Gary D. Libecap and Keith M. Shachat, 1990, An experimetric study of the Cournot theory of firm behavior, manuscript, University of Arizona.

Cournot, Augustin, 1960, orig. 1838, N.T. Bacon trans., Researches into the mathematical principles of the theory of wealth, (Kelly, New York).

Davidson, Carl and Raymond Deneckere, 1986, Long-run competition in capacity, short-run competition in price, and the cournot model, RAND Journal of Economics 17, 1986.

Davis, Douglas D., and Charles A. Holt, 1993, Experimental economics (Princeton University Press, Princeton).

Fisher, Franklin, 1961, The stability of the Cournot oligopoly solution: the effects of speeds of adjustment and increasing marginal costs, Review of Economic Studies 29, 125-135

Fouraker, Lawrence E., and Sidney Siegel, 1963, Bargaining behavior (McGraw-Hill, New York).

Herk, Leonard F., 1993, Consumer choice and Cournot behavior in capacity-constrained duopoly competition, RAND Journal of Economics 24, 399-417.

Holt, Charles A., 1985, An experimental test of the consistent-conjectures hypothesis, American Economic Review 75, 462-470.

Holt, Charles A., 1995, Industrial organization: a survey of laboratory research, in Alvin Roth and John Kagel, eds., Handbook of Experimental Economics (Princeton University Press, Princeton).

Kementa, Jan, 1978 Elements of Econometrics (Macmillian, New York).

Ketcham, Jon, Vernon L. Smith, and Arlington W. Williams, 1984, A comparison of posted-offer and double-auction pricing institutions, Review of Economic Studies 51, 595-614.

Kreps, David M. and Jose A. Scheinkman, 1983, Quantity precommitment and Bertrand competition yield Cournot outcomes, Bell Journal of Economics 14, 326-337.

Kruse, Jamie B., 1993, Nash equilibrium and buyer rationing rules: experimental evidence, Economic Inquiry 31, 631-646.

Kruse, Jamie, Steve Rassenti, Stan S. Reynolds and Vernon L. Smith, 1994, Bertrand-Edgeworth competition in experimental markets, Econometrica 62, 343-371.

Mason, Charles F., Owen Phillips, and Douglas B. Redington, 1991, The role of gender in a non-cooperative game, Journal of Economic Behavior and Organization 15, 215-235.

Mestelman, Stuart, Deborah Welland and J. Douglas Welland, 1987, Advance production in posted-offer markets, Journal of Economic Behavior and Organization 8, 249-264.

Mestelman, Stuart and J. Douglas Welland, 1988, Advance production in experimental markets, Review of Economic Studies 55, 641-654.

Noussair, Charles N, Charles R. Plott and Raymond G. Riezman, 1995, An experimental investigation of the patterns of international trade, American Economic Review 85, 462-491.

Pearce, Douglas K., 1984 Rationalizable strategic behavior and the problem of perfection. Econometrica 52, 1029-1050.

Phillips, Owen R., and Charles F. Mason, 1992, Mutual forbearance in experimental conglomerate markets, RAND Journal of Economics 23, 395-414.

Plott, Charles R., 1989, An updated review of industrial organization applications of experimental methods," in Richard Schmalensee and Robert Willig, eds., Handbook of Industrial Organization, vol. 2. (North Holland, Amsterdam) 1109-1176.

Plott, Charles R., Alexandre Borges Sugiyama, and Gilad Elbaz, 1995, Economics of scale, natural monopoly, and imperfect competition in an experimental market, Southern Economic Journal 61, 261-287

Rassenti, Stephen, Stanley S. Reynolds and Vernon L. Smith, 1996, Adaptation and convergence of behavior in repeated experimental Cournot games, manuscript, University of Arizona.

Sherman, Roger, 1972, Oligopoly, an empirical approach (Lexington Books, Lexington).

Theocharis, R. D., 1960, On the stability of the Cournot solution on the oligopoly problem, Review of Economic Studies 27, 133-134.

Vega-Redondo, Fernando, 1996, The evolution of Walrasian behavior, manuscript, Universidad de Alicante.

Wellford, Charissa P., 1990, Horizontal mergers: policy and performance, Ph. D. dissertation, University of Arizona.

White, Halbert,.1980, A heteroskedasticity-consistent covariance matrix estimator and direct test for heteroskedasticity, Econometrica, 48, 817-838.

Table 1. Convergence Patterns of Mean Transactions Prices, Sales Quantities, and Earnings

Posted Offer Sessions

	(1) Prices			(2) Sales Qtys.			(3) Earnings	
$\beta_{po\,1}$	207.15^{w}	(12.20)	$\beta_{po\,1}$	5.52^{w}	(0.63)	$\beta_{po\,1}$	348.26^{w}	(30.97)
$\beta_{po\,2}$	103.21^{*}	(31.53)	$\beta_{po\,2}$	5.38^{w}	(0.63)	$\beta_{po\,2}$	179.32^{*}	(36.70)
$\beta_{po\,3}$	118.01^{w}	(2.25)	$\beta_{po\,3}$	6.92^{w}	(0.85)	$\beta_{po\,3}$	159.44^{*}	(6.83)
$\beta_{po\,4}$	172.51^{w}	(43.22)	$\beta_{po\,4}$	6.92^{w}	(2.41)	$\beta_{po\,4}$	148.70^{*}	(88.13)
$\beta_{po\,5}$	165.17^{w}	(5.37)	$\beta_{po\,5}$	6.68^{w}	(0.89)	$\beta_{po\,5}$	324.04^{w}	(18.86)
$\beta_{po\,6}$	88.18^{*}	(21.79)	$\beta_{po\,6}$	7.85^{w}	(0.29)	$\beta_{po\,6}$	139.13^{*}	(46.31)
$\beta_{po\,fin}$	79.51^{s}	(0.97)	$\beta_{po\,fin}$	11.47	(0.07)	$\beta_{po\,fin}$	141.01^{s}	(2.90)
P_e	80.00		Q_e	12		π_e	140	
$\bar{R}^2$	0.95		$\bar{R}^2$	0.95		$\bar{R}^2$	0.82	
n	234		n	234		n	234	

Advance Production Sessions

	(1) Prices			(2) Sales Qtys.			(3) Earnings	
$\beta_{ap\,1}$	95.72^{*}	(4.51)	$\beta_{ap\,1}$	9.14	(0.33)	$\beta_{ap\,1}$	42.20^{w}	(16.87)
$\beta_{ap\,2}$	132.90	(3.54)	$\beta_{ap\,2}$	6.16	(0.88)	$\beta_{ap\,2}$	130.85^{w}	(13.43)
$\beta_{ap\,3}$	150.64	(21.87)	$\beta_{ap\,3}$	6.65	(0.59)	$\beta_{ap\,3}$	47.23^{w}	(88.12)
$\beta_{ap\,4}$	188.20	(6.37)	$\beta_{ap\,4}$	4.72^{w}	(0.87)	$\beta_{ap\,4}$	82.83^{w}	(39.90)
$\beta_{ap\,5}$	128.19	(3.40)	$\beta_{ap\,5}$	7.43	(0.36)	$\beta_{ap\,5}$	80.90^{w}	(23.83)
$\beta_{ap\,6}$	104.58^{*}	(7.23)	$\beta_{ap\,6}$	6.64	(0.63)	$\beta_{ap\,6}$	56.39^{w}	(34.05)
$\beta_{ap\,fin}$	95.44_{no}	(2.01)	$\beta_{ap\,fin}$	10.72_{no}	(0.11)	$\beta_{ap\,fin}$	183.18_{no}	(6.66)
P_c	140		Q_c	9		π_c	300	
$\bar{R}^2$	0.51		$\bar{R}^2$	0.60		$\bar{R}^2$	0.13	
n	217		n	217		n	217	

Note: Coefficient (standard errors). Estimates are corrected for first order autocorrelation. White's method is used to correct standard errors for heteroskedasticity. "*w*" indicates that a series passes weak convergence, "*" indicates that the weak convergence test does not apply. "*s* "indicates that the treatment passes strong convergence. "*no*" indicates no overlap in 95% confidence intervals about $\beta_{ap\,fin}$ and $\beta_{ap\,fin}$ estimates for a series.

Table 2. Individual Decisions and Optimality

Period Segment:	1-5	6-10	11-15	16-20	21-25	26-30	31-35
(1) Sizable deviations from Optimal Price: % of total relevant price choices.*	78.8	54.0	34.6	31.1	20.8	15.3	15.5
(2) Deviation from Best Response (Positive Dev.): % of total quantity choices.	95.5 (94.4)	92.2 (88.9)	88.9 (88.9)	77.8 (73.3)	81.1 (77.8)	80.0 (65.6)	82.1 (70.3)
(3) Deviation from Rationalizability: % of total quantity choices.	38.9	27.8	18.9	13.3	17.8	6.7	4.4

* Price postings are classified as sizable deviations if they are either above the market clearing price, or less than two cents below the market-clearing price. The set of relevant price choices consists of all decisions where posting the market clearing price is a dominant strategy.

Table 3. Nash Equilibria Assuming S1, S2 and S3 Maximize Earnings

Earnings for **S1**, S3 Quantity = 1

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	2.55	4.30	5.45	6.20	6.40	6.20
2	2.35	3.90	4.85	5.40	5.40	5.00
3	2.15	3.50	4.25	4.60	4.40	3.80
4	1.95	3.10	3.65	3.80	3.40	2.60
5	1.75	2.70	3.05	3.00	2.40	1.40
6	1.55	2.30	2.45	2.20	1.40	0.20

Earnings for **S2**, S3 Quantity = 1

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	2.55	2.35	2.15	1.95	1.75	1.55
2	4.30	3.90	3.50	3.10	2.70	2.30
3	5.45	4.85	4.25	3.65	3.05	2.45
4	6.20	5.40	4.60	3.80	3.00	2.20
5	6.40	5.40	4.40	3.40	2.40	1.40
6	6.20	5.00	3.80	2.60	1.40	0.20

Earnings for **S3**, S3 Quantity = 1

S2 Qty \ S1 Qty	1	2	3	4	5	6
1	2.55	2.35	2.15	1.95	1.75	1.55
2	2.35	2.15	1.95	1.75	1.55	1.35
3	2.15	1.95	1.75	1.55	1.35	1.15
4	1.95	1.75	1.55	1.35	1.15	0.95
5	1.75	1.55	1.35	1.15	0.95	0.75
6	1.55	1.35	1.15	0.95	0.75	0.55

Earnings for **S1**, S3 Quantity = 2

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	2.35	3.90	4.85	5.40	5.40	5.00
2	2.15	3.50	4.25	4.60	4.40	3.80
3	1.95	3.10	3.65	3.80	3.40	2.60
4	1.75	2.70	3.05	3.00	2.40	1.40
5	1.55	2.30	2.45	2.20	1.40	0.20
6	1.35	1.90	1.85	1.40	0.40	-1.0

Earnings for **S2**, S3 Quantity = 2

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	2.35	2.15	1.95	1.75	1.55	1.35
2	3.90	3.50	3.10	2.70	2.30	1.90
3	4.85	4.25	3.65	3.05	2.45	1.85
4	5.40	4.60	3.80	3.00	2.20	1.40
5	5.40	4.40	3.40	2.40	1.40	0.40
6	5.00	3.80	2.60	1.40	0.20	-1.0

Earnings for **S3**, S3 Quantity = 2

S2 Qty \ S1 Qty	1	2	3	4	5	6
1	4.30	3.90	3.50	3.10	2.70	2.30
2	3.90	3.50	3.10	2.70	2.30	1.90
3	3.50	3.10	2.70	2.30	1.90	1.50
4	3.10	2.70	2.30	1.90	1.50	1.10
5	2.70	2.30	1.90	1.50	1.10	0.70
6	2.30	1.90	1.50	1.10	0.70	0.30

Table 4. Nash Equilibria Assuming S1, S2 and S3 Maximize Earnings

Earnings for **S1**, S3 Quantity = 3

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	2.15	3.50	4.25	4.60	4.40	3.80
2	1.95	3.10	3.65	3.80	3.40	2.60
3	1.75	2.70	3.05	3.00	2.40	1.40
4	1.55	2.30	2.45	2.20	1.40	0.20
5	1.35	1.90	1.85	1.40	0.40	-1.0
6	1.15	1.50	1.25	0.60	-.60	-2.2

Earnings for **S2**, S3 Quantity = 3

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	2.15	1.95	1.75	1.55	1.35	1.15
2	3.50	3.10	2.70	2.30	1.90	1.50
3	4.25	3.65	3.05	2.45	1.85	1.25
4	4.60	3.80	3.00	2.20	1.40	0.60
5	4.40	3.40	2.40	1.40	0.40	-.60
6	3.80	2.60	1.40	0.20	-1.0	-2.2

Earnings for **S3**, S3 Quantity = 3

S2 Qty \ S1 Qty	1	2	3	4	5	6
1	5.45	4.85	4.25	3.65	3.05	2.45
2	4.85	4.25	3.65	3.05	2.45	1.85
3	4.25	3.65	**3.05**	2.45	1.85	1.25
4	3.65	3.05	2.45	1.85	1.25	0.65
5	3.05	2.45	1.85	1.25	0.65	0.05
6	2.45	1.85	1.25	0.65	0.05	-.55

Earnings for **S1**, S3 Quantity = 4

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	1.95	3.10	3.65	3.80	3.40	2.60
2	1.75	2.70	3.05	3.00	2.40	1.40
3	1.55	2.30	2.45	2.20	1.40	0.20
4	1.35	1.90	1.85	1.40	0.40	-1.0
5	1.15	1.50	1.25	0.60	-.60	-2.2
6	0.95	1.10	0.65	-.20	-1.6	-3.4

Earnings for **S2**, S3 Quantity = 4

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	1.95	1.75	1.55	1.35	1.15	0.95
2	3.10	2.70	2.30	1.90	1.50	1.10
3	3.65	3.05	2.45	1.85	1.25	0.65
4	3.80	3.00	2.20	1.40	0.60	-.20
5	3.40	2.40	1.40	0.40	-.60	-1.6
6	2.60	1.40	0.20	-1.0	-2.2	-3.4

Earnings for **S3**, S3 Quantity = 4

S2 Qty \ S1 Qty	1	2	3	4	5	6
1	6.20	5.40	4.60	3.80	3.00	2.20
2	5.40	4.60	3.80	3.00	2.20	1.40
3	4.60	3.80	3.00	2.20	1.40	0.60
4	3.80	3.00	2.20	1.40	0.60	-.20
5	3.00	2.20	1.40	0.60	-.20	-1.0
6	2.20	1.40	0.60	-.20	-1.0	-1.8

Table 5. Nash Equilibria Assuming S1, S2 and S3 Maximize Earnings

Earnings for **S1**, S3 Quantity = 5

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	1.75	2.70	3.05	3.00	2.40	1.4
2	1.55	2.30	2.45	2.20	1.40	0.20
3	1.35	1.90	1.85	1.40	0.40	-1.0
4	1.15	1.50	1.25	0.60	-.60	-2.2
5	0.95	1.10	0.65	-.20	-1.6	-3.4
6	0.75	0.70	0.05	-1.0	-2.6	-3.4

Earnings for **S2**, S3 Quantity = 5

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	1.75	1.55	1.35	1.15	0.95	0.75
2	2.70	2.30	1.90	1.50	1.10	0.70
3	3.05	2.45	1.85	1.25	0.65	0.05
4	3.00	2.20	1.40	0.60	-.20	-1.0
5	2.40	1.40	0.40	-.60	-1.6	-2.6
6	1.40	0.20	-1.0	-2.2	-3.4	-3.4

Earnings for **S3**, S3 Quantity = 5

S2 Qty \ S1 Qty	1	2	3	4	5	6
1	6.40	5.40	4.40	3.40	2.40	1.40
2	5.40	4.40	3.40	2.40	1.40	0.40
3	4.40	3.40	2.40	1.40	0.40	-.60
4	3.40	2.40	1.40	0.40	-.60	-1.6
5	2.40	1.40	0.40	-.60	-1.6	-2.6
6	1.40	0.40	-.60	-1.6	-2.6	-2.6

Earnings for **S1**, S3 Quantity = 6

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	1.55	2.30	2.45	2.20	2.40	1.40
2	1.35	1.90	1.85	1.40	1.40	0.20
3	1.15	1.50	1.25	0.60	0.40	-1.0
4	0.95	1.10	0.65	-.20	-.60	-2.2
5	0.75	0.70	0.05	-1.0	-1.6	-3.4
6	0.55	0.30	-.55	-1.8	-2.6	-3.4

Earnings for **S2**, S3 Quantity = 6

S2 Qty \ S1 Quantity	1	2	3	4	5	6
1	1.55	1.35	1.15	0.95	0.75	0.55
2	2.30	1.90	1.50	1.10	0.70	0.30
3	2.45	1.85	1.25	0.65	0.05	-.55
4	2.20	1.40	0.60	-.20	-1.0	-1.8
5	1.40	0.40	-.60	-1.6	-2.6	-2.6
6	0.20	-1.0	-2.2	-3.4	-3.4	-3.4

Earnings for **S3**, S3 Quantity = 6

S2 Qty \ S1 Qty	1	2	3	4	5	6
1	6.20	5.00	3.80	2.60	1.40	0.20
2	5.00	3.80	2.60	1.40	0.20	-1.0
3	3.80	2.60	1.40	0.20	-1.0	-2.2
4	2.60	1.40	0.20	-1.0	-2.2	-3.4
5	1.40	0.20	-1.0	-2.2	-3.4	-3.4
6	0.20	-1.0	-2.2	-3.4	-3.4	-3.4

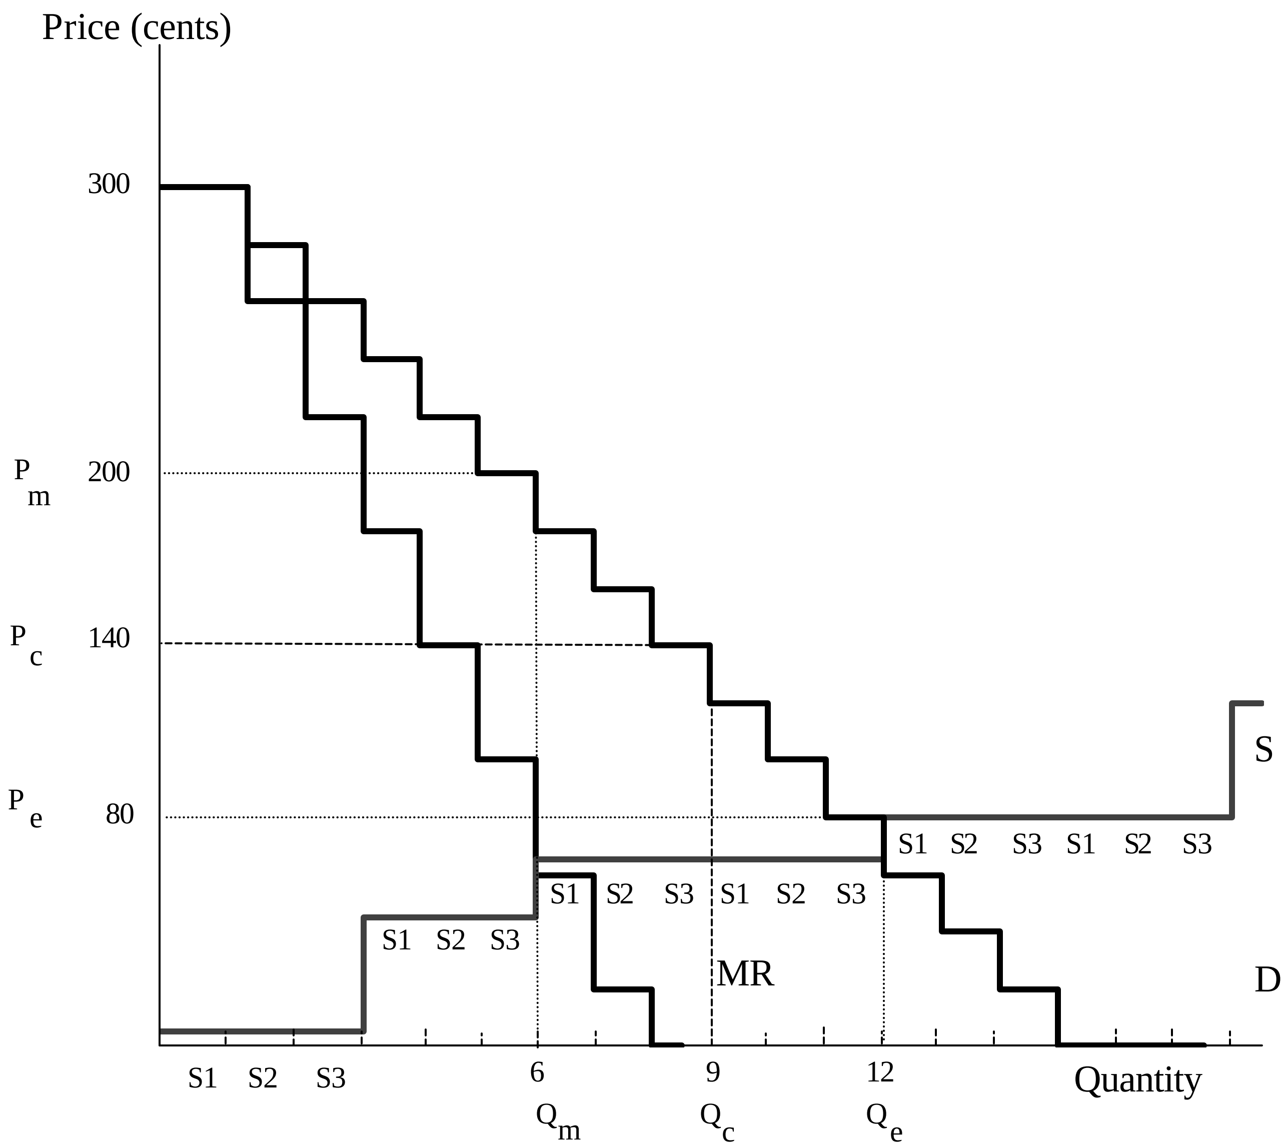

Figure 1. Aggregate Supply and Demand Arrays.

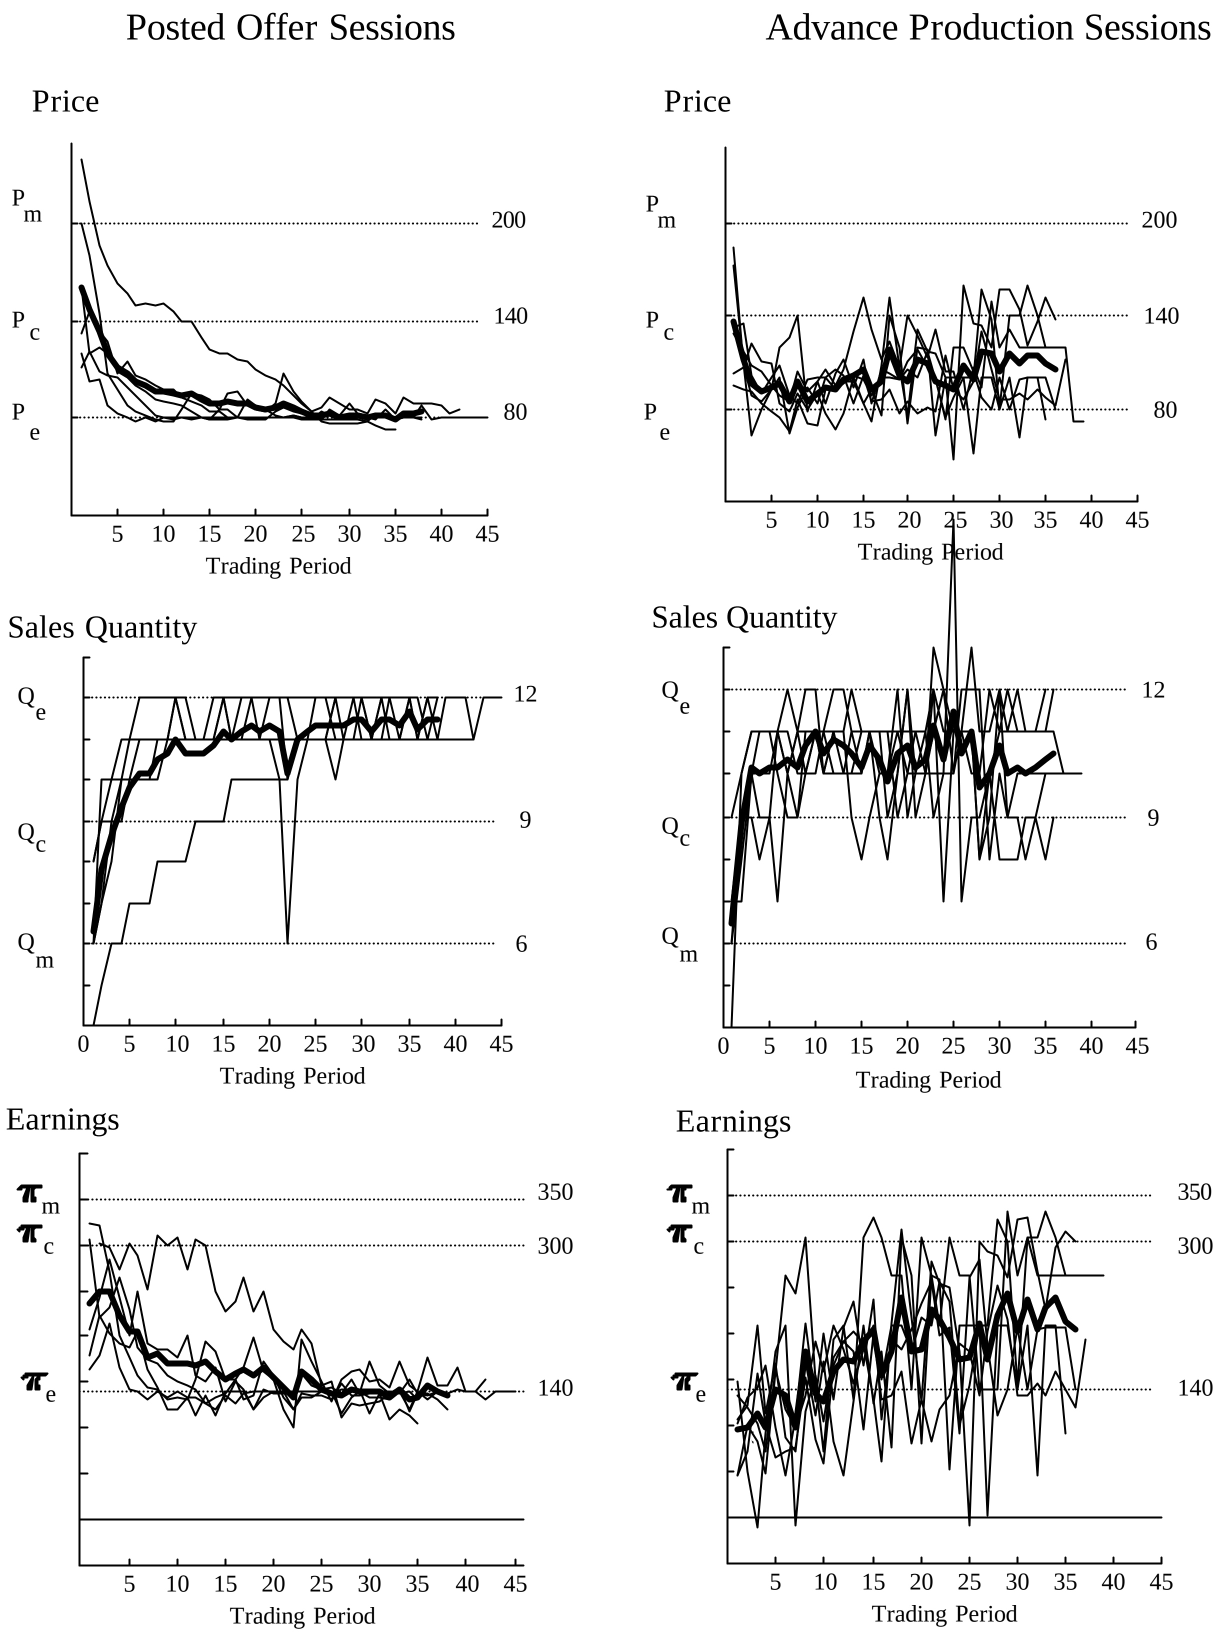

Figure 3. Mean Transactions Prices, Quantities, and Mean Earnings Paths for Posted Offer (PO) and Advance Production (AP) treatments. Key: In each panel, the thin lines illustrate series for the six sessions in each treatment cell. The thick line illustrates the average series value for the treatment.

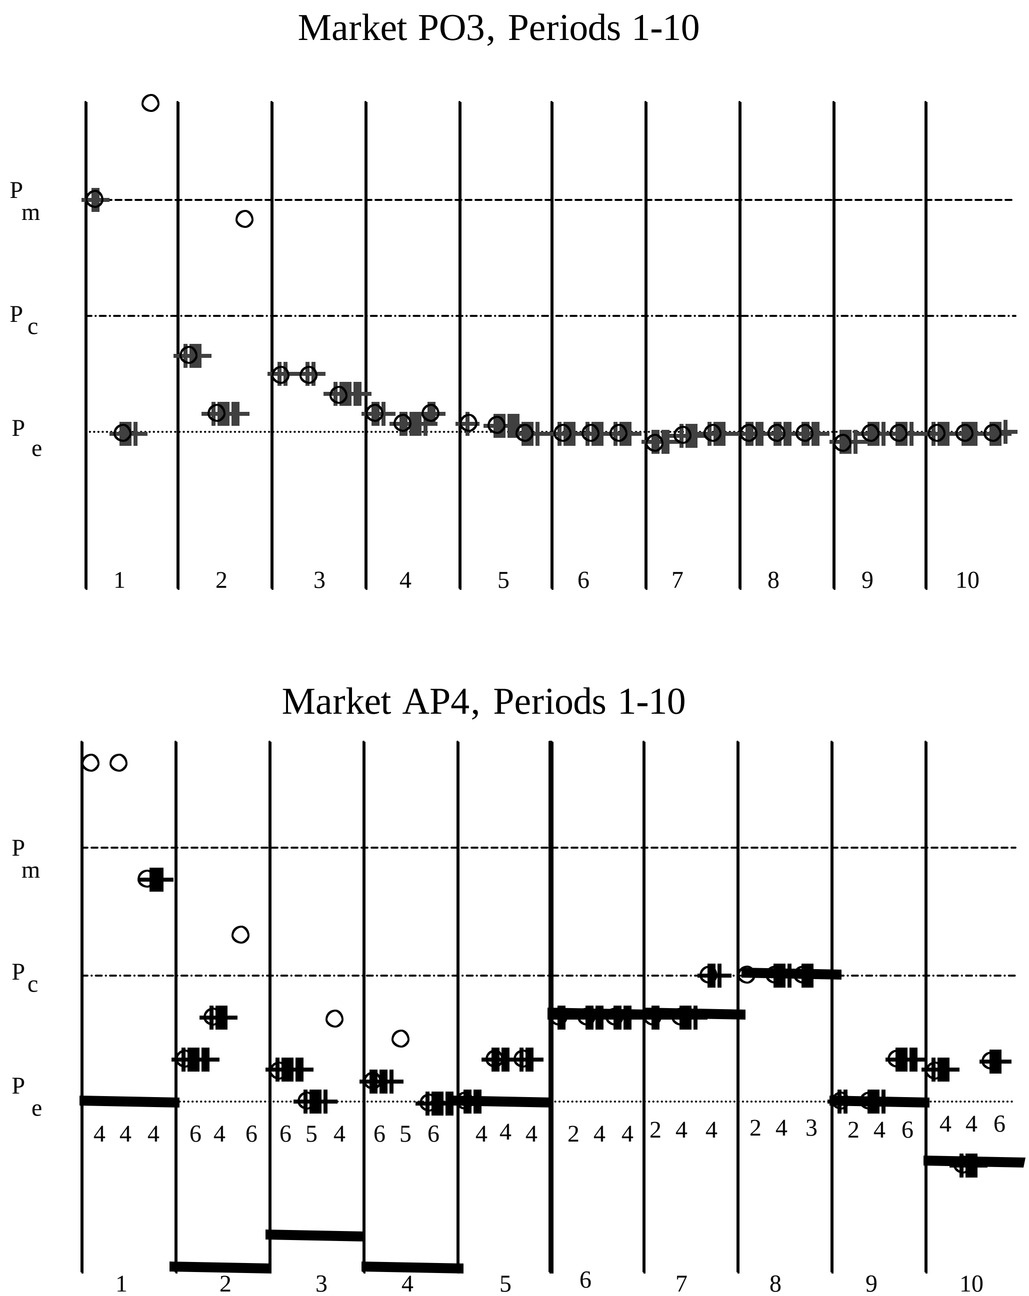

Figure 2. Contract Sequences for the First 10 Trading Periods of a Representative Posted Offer (PO) Market, and a representative Advance Production (AP) Market. Key: Within each panel trading periods are illustrated as vertical stripes. Dots in the stripes represent the prices posted by sellers S1, S2 and S3, respectively, and crosses extending from the dots represent contracts. In the lower (AP) panel, offer quantities for S1, S2 and S3 are printed in respective order below the dotted line illustrating the competitive price Pe. The bold horizontal segments in this panel illustrate the market clearing price, given aggregate production for the period.