Methodological Issues in Analyzing Market Dynamics.

Barcelona Workshop
June, 2014.

by

Ariel Pakes
(Harvard University and the NBER).

Background: Methodological Developments in IO.

- We have been developing tools that enable us to better analyze market outcomes.

- Common thread: emphasis on incorporating the institutional detail that is needed to make sense of the data used in analyzing (i) the likely causes of historical events, or (ii) the likely responses to environmental and policy changes.

- Focus. **Incorporate**
(i) **heterogeneity** (in plant productivity, products demanded, bidders and/or consumers) and where possible,
(ii) **equilibrium conditions** (Nash in prices or quantities, and extensions designed to analyze allocations in network, platform, and vertical markets).

We largely relied on earlier work by our game theory colleagues for the analytic frameworks.

- Each agent's actions affect all agents' pay-offs, and

- At the "equilibrium" or "rest point"
(i) agents have correct perceptions, and
(ii) the system is in some form of Nash equilibrium (policies such that no agent has an incentive to deviate).

- Our contribution is the development of an ability to adapt the analysis to the richness of different real world institutions.

This is the second time I have had the opportunity to revisit progress. In the first, nine years ago, I focused on on the developments in "static" analysis. Now I think we can make the following claim.

Claim 1: The tools developed for the analysis of market allocations conditional on the "state variables" of the problem pass a "market test" for success. I say this because: (i) they have been incorporated into applied work in many fields, and (ii) they fit surprsingly well (at least in the cross-section).

Market use test.

- Tools that we modified to incorporate more realism have seen extensive independent use to analyze productivity and demand in many economic studies.

- Full equilibrium allocation models are often used to analyze market allocatons when price taking or quantity setting is a reasonable approximation.

- We have developed procedures for analyzing equilibrium allocations in more complex market situations and they are being continually improved. E.g.'s auction markets, and importantly in vertical markets characterized by multi-agent bargainng with externalities.

Indeed there are by now many instance where the methodology has been used up by public agencies, consultancies, and, to some extent, by private firms.

E.g. of Fit: Nash Pricing Equation. Tom Wollman's dissertation (in progress) on commercial trucks. Estimate BLP demand, regress implied Nash markup on instruments to get $\widehat{markup}$ (R^2=.44 or .46 with time dummies; low relative to more sophisticated IVs).
Look at fit & is coefficient of $\widehat{markup} \approx 1$?

Table 1: Fit of Pricing Equilibrium.

	Price	(S.E.)	Price	(S.E.)
Gross Weight	.36	(0.01)	.36	(.003)
Cab-over	.13	(0.01)	.13	(0.01)
Compact front	-.19	(0.04)	0.21	(0.03)
long cab	-.01	(0.04)	0.03	(0.03)
Wage	.08	(.003)	0.08	(.003)
$\widehat{Markup}$	.92	(0.31)	1.12	(0.22)
Time dummies?	No	n.r.	Yes	n.r.
R^2	0.86	n.r.	0.94	n.r.

Nobs=1,777; firms=16; t=1992-2012; Heter-cons s.e.

What About “Dynamics”? Use textbook distinction: (i) static models solve for profits conditional on state variables, (ii) dynamics analyzes the evolution of those state variables.

Claim 2: Our methods of analyzing the evolution of the “state variables” would not get nearly as high a grade on a market test.

• We have made significant progress on methodological issues in dynamic analysis, and there have been some exceptionally good empirical work on the dynamics of a few industries.

• However our modes of dynamic analysis of market outcomes have not passed beyond the confines of a few diligent and creative researchers.

I want to spend the rest of this talk asking;
(i) why, and
(ii) what types of research might move us forward with empirical work on dynamics.

The discussion will be at a relatively "high level", starting with an overview of the conclusions I will reach with references to details where needed. I then come back to the thought process which led to them. **Apologies** for no references to important papers whose conclusions I do not use directly.

Conclusions.

- We have relied too much on overly demanding notions of equilibrium that theorists have used in guiding their analytic studies of market dynamics.

- At least for markets which can reasonably be approximated by transitions to states which have been visited repeatedly, alternative equilibrium notions that demand less of both the analyst and economic agents;
  (i) are available,
  (ii) likely to better approximate behavior,
  (iii) and can be analyzed with (perturbations of) many of the tools that we have developed to analyze the more demading equilibrium notions.

- Finally, more empirical work is needed on;
  (i) how firms learn to deal with changes in their environment,
  (ii) how we can integrate such changes into frameworks which feature "common sense" notions of equilibrium on states that are visited repeatedly,
  (iii) what additional (to the common sense) properties can we expect firm behavior to satisfy when on these recurrent states .

How did empirical dynamics develop? As in static analysis we took our frameworks largely from theory and tried to incorporate the institutions that seemed necessary to analyze actual market settings. The relevant theory

- made assumptions which implied that the state variables evolved as a Markov process

- and imposed that the equililbrium satisfies some form of Markov Perfection (agents perceptions are correct and no agent has an incentive to deviate at any value of the state variables).

When we added the relevant institutional background, the dynamic analysis became quite complex. This became evident when we tried to use the Markov Perfect notions to structure

- the estimation of parameters, or to

- compute the fixed points that defined the equilibria or rest points of the system.

Our response. Keep the equilibrium notion and develop techniques to make it easier to circumvent the estimation and computational problems. Useful contribution in this regard:

- The development of estimation techniques that circumvent the problem of repeatedly computing equilibria (that do not require a nested fixed point algorithm).

- the use of approximations and functional forms for primitives which enabled us to compute equilibria quicker and/or with less memory requirements.

However even after integrating the new dynamic tools, applied research only had limited ability to accomodate the relevant institutions. As a result;
(i) many dynamic issues are thought to be too difficult to analyze in a realistic way and hence are under analyzed (e.g; empirical work on the dynamics of mergers); and
(ii) even where these tools were used in empirical work the specifications analyzed were dictated as much by computational as by institutional considerations.

I want to take a step back and consider how we might rectify the situation. I start with the following (non-novel) observation.

The fact that the Markov Perfect framework becomes unwieldily when confronted by the complexity of real world institutions, not only

- limits our ability to do empirical analysis of market dynamics,

- it also raises the question of whether some other notion of equilibria will better approximate agents' behavior.

This raises the following question. If we abandon Markov Perfection can we both

- better approximate agents' behavior and,

- enlarge the set of dynamic questions we are able to answer.

First question: what underlies the complexity of our Markov Perfect models? When we try to incorporate the insitutional background we think is essential to the analysis we find:

1. That the agent is required to access a large amount of information (all state variables) and either compute, or learn and retain, an unrealistic number of strategies (one for each information set); moreover

2. Direct computation of these strategies is a challenging task, often too challenging to think an agent could accomplish it.

What determines the cardinality of the information sets "needed" for symmetric information empirical work? The state variables should include the major determinants of: demand, cost, and any additional historical information that effects equilibrium strategies.

- To see how demanding this can be, consider markets where consumer, as well as producer, choices are dynamic (e.g.'s; durable, experience, or network goods); need the distribution of; current stocks $\times$ household characteristics, production costs, ...

• In a symmetric information MPE an agent would have to access all state variables, and then either compute, or learn and retain, policies from each distinct information set. To compute policies we would need to compute a doubly nested fixed point.

Obvious Theory Fix: Assume agents only have access to a subset of the state variables.

- Since agents presumably know their own characteristics and these tend to be persistent, we would need to allow for assymetric information: if we maintained "perfectness" theory would lead us to a "Bayesian" Markov Perfect solution.

Bayesian Perfect has a problem.

- The computational burden of computing its strategies makes it a **non-starter** for even the simplest (realistic) applied problem. The additional burden results from the need to compute posteriors, as well as optimal policies; and the requirement that they be consistent with one another.

This does not mean that the policies can not be obtained in other ways. An obvious possibility is that they are learned from combining data on past behavior and market outcomes.

- They would have to learn about;

(i) primitives (some empirical work on this),
(ii) the likely behavior of their competitors, and
(iii) market outcomes given primitives, competitor behavior, and their own policies.

• There is surprising little empirical evidence on how firms formulate their perceptions about either other firms' behavior, or on the impact of their own strategies given primitives and the actions of competitors.

• So I am going to turn to an example (by U. Doraszelski, G. Lewis and myself). We study the bids from the date the British Electric Utility market for frequency response opened.

Background. Frequency response is a product required by the system operator in electricity markets to keep the system running smoothly (it brings demand into balance with supply and keeps system frequency in a particular "hertz" band).

A private firm runs the grid and buys frequency response from electricity generating firms. Prior to November 2005, generators in the UK were required to provide frequency response to the system operator at a fixed system-wide price. Then the market was deregulated and generators were allowed to bid into an auction market, setting the stage for price competition.

• On the one hand this is a simple environment in which to study a new market
(i) Firms had been regulated to provide frequency response prior to the start of the auction market and so likely knew demand and supply primitives which vary with season.
(ii) The game is repeated monthly, and there is no evidence of collusion.
(iii) Eight players have 75-80% of total revenue. One (Drax) has 1/4.

• On the other; we study a whole new insitution for aggregating firm decisions into market outcomes. In particular initially the firms in this market did not have a history of bids and outcomes to rely on. Most often we are studying smaller changes with more history available on policies and outcomes in similar states.

Graph on bids of largest 8 players.

- Firms start out with low bids (about equal to regulated reimbursement), and at different times begin experimenting (not in a co-ordinated way).

- After about 18 periods, they all trend downward.

- Graph on cross-sectional variance of bids. The variance in the share weighted average of bids falls be a factor of 5 to 10 and stays low. Generated electricity is not quite (but almost) a homogenous good (depends on location and reliability of generator).

- After about 42 periods they seem to have settled about a bid (with small perturbations period to period), and stay that way until about 2010.

- Early 2010 national grid starts using long term contracts with Drax. It takes a while to readjust, but not nearly as long or as wild a perturbation as we see at the beginning of the market.

- Indeed, one sense in which the graph is encouraging: **firms bids evenutally stabilize or "converge"**.

- The initial convergence process was quite complicated, even after the initial period of experimentation. The yellow line is Drax. During the convergence process Drax twice tries to raise price, and when no-one follows, it drops back down.

- The smaller change (Drax's contracts) induced shorter convergence times and less variant bids before convergence, but still generated some adjustments before convergence. This may be more characteristic of the changes we see in existing markets.

Possible Takeaways: Theory & Example.

• In many settings we study there are likely to be persistent differences in informations sets.

• It is unlikely (if not impossible) that agents obtain their policies by computing a Bayesian Perfect equilibrium.

• In markets that are reasonably stable, we might treat firms as having learned a set of strategies; at least for information sets visited repeatedly.

• When institutions change it takes some time for firms' policies to stabilize. The experimentation in the interim may be to find out competitor's reactions as well as to explore outcomes given primitives and competitor's play.

• The time required and the extent of variability in the interim is likely to be directly related to the magnitude of the change.

How Might We Proceed?

• It might be useful to start by dividing the set of issues to be analyzed into two subsets;
(i) those for a learning period when policies are "transient", and
(ii) those for the period when we observe "recurrent" strategies.

The analysis of the recurrent case seems easier, but still opens up at least one important question.

- Easier because the strategies likely satisfy a Nash condition of some sort; else someone has an incentive to deviate.

- The question that remains is the form of the Nash Condition.

Assertion. When working with recurrent policies the process those policies generate will be well approximated by a time homogenous Markov process (of finite order).

• Theory may lead you to worry about this assertion since dynamic games of assymetric information may generate an infinite regress. I regard this largely as a "Red Herring". Information access and retention conditions, as well as approximation theory, are likely to trump it.

• Still, it would be good to have a formal argument which told us what conditions we need for approximations to be convergent, as that would give us some idea of when their might be a problem.

“Transient Policies”.

There are many theoretical options here, and little empirical work to guide us in sorting them out. I have much less to say here, though there is some past work.

• Firms need to learn about three phenomena:
(i) primitives and their evolution (there is empirical work here; e.g. learning by doing, learning from others,...)
(ii) the likely actions of competitors (mostly theory and lab experiments), &
(iii) in new markets, the allocational response to a change in my policy given primitives and the actions of my competitors (e.g. in our auction, the response to my bid, conditional on demand and the bids of my competitors).

• Just as we allow for heterogeneity in our static equilibrium models we should be cognizant of the fact that learning may occur in different ways in different environments. The little empirical work I have seen indicates that

- Agents often differ both in: (i) initial priors and (ii) incentives to learn,

- Firms often give less weight to other firm's experiences then their own, and

- The best approximation to the learning process need not be Bayesian. Alternatives include;
  - "Statistical Learning" models that use various estimators,
  - "Reinforcement Learning" models which make it more likely to mimic behavior that was successful in the past (and more

generally, the OR or "machine" approximations to learning processes).

- These relatively simple learning algorithms are likely to be most suitable for changes in an existing (rather than a new) market. As a result they could be crucial for studying counterfactuals, helping to select out possible equilibria where relevant.

- Models which explicitly incorporate experimentation might well be needed in markets that faced dramatic changes or when dealing with research and/or exploration.

What Conditions Can We Assume for the Rest Point at States that are Visited Repeatedly?

• We have more to say here. We expect, and should integrate into our modelling, that

1. Agents perceive that they are doing the best they can at each of these points, and that

2. These perceptions are consistent with what they observe.

• If there are additional conditions known about the institutional environment (and their often are) we should impose them: say knowledge of the firm decision-making process, or of the value that would be generated by counterfactual behavior. I come back to this below, and show that it can be used to limit the number of possible equilibria.

Formalizing Equilibrium Conditions.

• Denote the information set of firm i in period t by $J_{i,t}$. $J_{i,t}$ will contain both public and private information. Policies, say $m_{i,t} \in \mathcal{M}$, will be functions of $J_{i,t}$. For simplicity we will assume $\#\mathcal{M}$ is finite.

• Assume $J_{i,t}$ evolves as a (controlled) finite state Markov process on $\mathcal{J}$, (or at least can be adequately approximated by one); and that there are a finite number of firms ever active at one time. So a "state" of the system, is

$$s_t = \{J_{1,t}, \ldots, J_{n_t,t}\}.$$

• This implies that the points that are observed repeatedly define a recurrent subset of $\mathcal{S}$, say $\mathcal{R} \subset \mathcal{S}$, and that the policies at those points insure that s_t will wander into $\mathcal{R}$ in finite time, and after that $s_{t+\tau} \in \mathcal{R}$ w.p.1 forever.

Implications of Our Assumptions.

Let the agent's perception of the expected discounted value of current and future net cash flow were it to chose m at state J_i, be

$$W(m|J_i), \ \ \forall m \in \mathcal{M} \ \ \& \ \ \forall J_i \in \mathcal{J},$$

and let is perception of expected profits be

$$\pi^E(m|J_i).$$

Then two implications of our assumptions are that if $m^*(J_i)$ is the policy chosen at J_i, then

$$A. \ \ W(m^*|J_i) \geq W(m|J_i), \ \ \forall m \in \mathcal{M} \ \ \& \ \ \forall J_i \in \mathcal{J}$$

and

$B.$ $\forall J_i$ which is a component of an $s \in \mathcal{R}$

$$W(m^*(J_i)|J_i) = \pi^E(m^*|J_i) + \beta \sum_{J_i^{'}} W(m^*(J_i^{'})|J_i^{'})p^{e}(J_{i}^{'}|J$$

and if $E[\pi(\cdot)|J_i, J_{-i}]$ is the DGP's average of agent i's profits given (J_i, J_{-i}) then

$$\pi^E(m^*|J_i) \equiv \sum_{J_{-i}} E[\pi(\cdot)|J_i, J_{-i}]p^e(J_{-i}|J_i),$$

where $p^e(\cdot)$ denotes an empirical (& hence estimable) probability which, $\forall J_i \in \mathcal{J}$, satisfy

$$\Big\{ p^e(J_i^{'}|J_i) \equiv \frac{p^e(J_i^{'}, J_i)}{p^e(J_i)} \Big\}_{J_i^{'}},$$

while

$$\Big\{ p^e(J_{-i}|J_i) \equiv \frac{p^e(J_{-i}, J_i)}{p^e(J_i)} \Big\}_{J_{-i}}. \ \spadesuit$$

"Experience Based Equilibrium"

These are the conditions of an EBE (Fershtman and Pakes, 2012). Bayesian Perfect satisfy them, but so do weaker notions. Notes.

• **Computation.** Both agents and analyst can compute equilibria without calculating posterior distributions. The computational algorithm can be interpreted as a learning algorithm that is easy for agents to use to find equilibrium policies provided they have access to enough information, but that access requires a number of assumptions ....

• The algorithm uses some of the computational advances made in earlier literature (stochastic approximation) and can be augmented to use functional form approximations in complex cases (the OR literature on "TD learning").

• The algorithm is: (1) iterative and, (2) "asynchronous" (i.e. updates one point at a time). If k indexes iteration memory has;
(i) estimates of $\{W(\cdot)\}$, say$\{W^k(\cdot|\cdot)\}$, and
(ii) a location $s^k = \{J_{1,k}, \ldots, J_{n_k,k}\}$.
Both are updated each iteration.
A. $\{W^k(\cdot|\cdot)\}$ determines currently optimal policies at s^k which are used to update the location. $\Rightarrow s^k$ eventually wanders into $\mathcal{R}$ and stays there. $\mathcal{R}$ need not grow in any particular way with the number of state variables.
B. Use of stochastic approximation then replaces the need to integrate over possible future states to update continuation values. The update is a weighted average of two terms:
(a) Last iterations estimate of W, say W^k,
(b) An estimate of the expected value obtained as the realized profits earned in the current iteration plus the current estimate of the continuation value from the updated location.
• A + B eliminate the two sources of the curse of dimensionality in computation.

• **Estimation.** The equilibirum condition

$$W(m^*|J_i) \geq W(m|J_i)$$

is an inequality which can be used to generate (set) estimators of parameters without ever computing the value function (as in developments for Markov Perfect estimation). We use a perturbation (or "Euler") condition.

Example for estimation:
• J contains both public and private information. Let J^1 have the same public, but differnt private, information then J^2.
• If a firm is at J^1 it knows it could have played $m^*(J^2)$ and its competitors would respond by playing *on the equilbrium path* from J^2.
• Given data on play on the recurrent class we could simulate a sample path from J^2.

- This simulated path only visits points in the recurrent class generated by the DGP. So it would intersect the sample path from the DGP at a random stopping time with probability one and from that time forward the two paths would generate the same profits.
- So the expected difference in discounted profits from the period of the deviation to the random stopping time, should, when evaluated at the true parameter vector, be positive.

• **Interior and Boundary Points.** Partition $\mathcal{R}$ into "interior" and "boundary" points. Points at which there are feasible strategies which can lead outside of $\mathcal{R}$ are boundary points. Interior points are points that can only transit to other points in $\mathcal{R}$ no matter which (feasible) policy is chosen.

• **Problem: Multiplicity.** Our conditions only insure that perceptions of outcomes are correct for equilibrium play. At a boundary $s \in \mathcal{R}$ agent's perceptions for off the equilibrium path play are not tied down by actual outcomes. As a result different perceptions can support a wide range of equilibria.

Dealing with Multiplicity.

- In any empirical appliction the data will rule out equilibria. m^* is observable, at least for states in $\mathcal{R}$, and this implies inequalities on $W(m|\cdot)$. With enough data $W(m^*|\cdot)$ will also be observable up to a mean zero error.

- Impose "a priori" restrictions on beliefs:
(i) $m_i^*(J_i)$ for $J_i \notin s \in \mathcal{R}$; e.g. bring in external information on perceptions of the value of outcomes outside of $\mathcal{R}$, or
(ii) $m_i \neq m_i^*$ for $J_i \in s \in \mathcal{R}$; e.g. if agents can learn what outcomes would be for an $m_i \neq m_i^*$ impose that the associated $W(m|J_i)$ are consistent with those outcomes.

- "Restricted EBE", formalizes (ii) for a set of situations for which it seems relevant. Whether or not we can use (i) seems to be application specific, but it is easy to add these restricions if relevant.

Restricted EBE. When might agents be able to learn the value of actions off the equilibrium path? For $m_i \neq m_i^*$, they will need to learn:
(i) expected profits,
(ii) the distribution of future states, and
(ii) the continuation values from states reached.

• A capital accumulation game is a game where the distribution of the characteristics of the products a firm markets and the cost of producing them depends only on the firm's own action and exogenous events.

• In capital accumulation games in which the firm knows its demand function conditional on observable actions of competitors (e.g. prices) and its own cost functions, firms can always learn the value of $W(m|J_i)$, $\forall m \in \mathcal{M}$ and J_i a component of an s which is an interior point of $\mathcal{R}$.

Simple Electric Utility Eg.

Two firms: each has a vector of generators.
Firm's decisions: bid or not each generator. If not bid, do maintenance or not.
ISO: sum bid functions, intersect with demand (varies by day of the week), pay a uniform price to accepted electricity.

- $\omega \in \Omega$. Cost of producing electricy on each firm's generators. Cost increases stochastically with use, but reverts to a starting value if the firm goes down for maintenance.

- $m_i \in M_i$. Vector of $m_{i,r} \in \{0, 1, 2\}$; $0 \Rightarrow$ shutdown without maintenance, $1 \Rightarrow$ shutdown with maintenance, $2 \Rightarrow$ bid into market.

- $b(m_i) : m_i \to \{0, b_i\}^{n_i}$ where b_i is the fixed bid schedule of firm i. b observed. m not observed.

- d is demand on that day, f is maintenance cost ("investment"), $p = p(b(m_i), b(m_{-i}), d)$ is price, $q = q(b(m_i), b(m_{-i}), d)$ is allocated quantity vector, so realized profits are

$$\pi_{i,t} = \sum_r p_t q_{i,r,t} - \sum_r c_i(\omega_{i,r,t}, q_{i,r,t}) - f_i \sum_r \{m_{i,r,t} = 1\}$$

$$\equiv \pi_i(\omega_i, m_i, b(m_{-i}), d)$$

$m_{i,r,t} = 0 \;\Rightarrow \omega_{i,r,t+1} = \omega_{i,r,t}$,
$m_{i,r,t} = 1 \;\Rightarrow \omega_{i,r,t+1} = \overline{\omega}_{i,r}$ ($\overline{\omega}$=restart state),
$m_{i,r,t} = 2 \;\Rightarrow \; \omega_{i,r,t+1} = \omega_{i,r,t} - \eta_{i,r,t}$
$with\ P(\eta) > 0\ for\ \ \eta \in \{0, 1\}$.

Note $b(m)$ is the only signal sent in each period. $b(m_{-i,t-1})$ is a signal on $\omega_{-i,t-1}$ which is unobserved to i and is a determinant of $b(m_{-i,t})$ (and so $\pi_{i,t}$).

State of the game. $s_{i,t} = (J_{1,t}, \ldots J_{n_t,t}) \in \mathcal{S}$, and

$$J_{i,t} = (\xi_t, z_{i,t}) \in (\Omega(\xi), \Omega(z))$$

where

- $z_{i,t}$ represents private information (known only to i). Example: $\omega_{i,t} = z_{i,t}$.

- and ξ_t is public information (shared by all). Example $\xi_t = \{b(m_{1,\tau}), b(m_{2,\tau}), d_\tau\}_{\tau \le t}$.

Model Details.

Parameter	Firm B	Firm S
Number of Generators	2	3
Range of ω	0-4	0-4
MC @ $\omega = (0, 1, 2, 3)^*$	(20,60,80,100)	(50,100,130,170)
Capacity at Const MC	25	15
Costs of Maintenance	5,000	2,000

*MC is constant at this cost until capacity and then goes up linearly. At $\omega = 4$ the generator shuts down.

Firm S: small (gas fired) generators with high MC but low start up costs.
Firm B: large (coal fired) generators lower MC and higher start up costs.
Constant, small, elasticity of demand.

Computational Details.
• High initial conditions "insures" we try all strategies (induces a lot of experimentation).
• Convergence test is in terms of $\mathcal{L}^2(\mathcal{P}(\mathcal{R}))$ norm of percentage bias in estimates of W. 300 million iterations $\mathcal{L}^2(\mathcal{P}(\mathcal{R})) \approx .00005$.

The Economics of Alternative Environments: Planner vs AsI.

Base Case: Planner Strategy. Constrain planner to use the same bid function (compare just investment strategies). Never shuts down without doing maintenance. Weekdays: operates at almost full capacity. Maintenance done on weekend. Maintenance done about 15% of the periods for both B and S generators.

Base Case: AsI Equilibrium. Shuts down about 20% of the periods. However about half the time generators are shutdown they are not doing maintenance. Only does maintenance in about 10% of the periods. $\Rightarrow$ 25-30% *more* shutdown but 30% *less* maintenance than the social planner. Most (but not all) shutdown on weekends (just as social planner).

Base Case: Costs. Planner does more maintenance and can optimize maintenance jointly over large and small generators. $\Rightarrow$ much lower production costs and lower total costs per unit quantity.

• I.e. the planner produces more and has lower average total costs in a model in which marginal costs increase in quantity. Effect of increased maintenance.

Base Case: Prices and Quantities. Planners 2% more output on weekdays, with inelastic demand $\Rightarrow$ price fall of $\approx$ 10%.

• Planner's extra maintenance makes it optimal for it to bid in more and therefore keep price down, and it internalizes the extra CS. AsI firms do not.

• Even the social planner has weekday prices that are 20% higher than weekend prices (the AsI difference is larger). With these primitives large weekend/weekday price differences are "optimal".

	Base Case		
	SP	AsI	FI
Panel A: Strategies.			
Firm B: Shutdown and Maintenance.			
Shutdown %	14.52	19.96	12.31
Maintenance %	14.52	10.1	10.9
Firm S: Shutdown and Maintenance.			
Shutdown %	16.85	21.48	20.74
Maintenance %	16.85	9.83	9.91
Firm B: Operating Generators (by day).			
Saturday	1,41	1.08	1.72
Sunday	.88	1.21	1.65
Weekday Ave.	1.93	1.78	1.78
Firm S: Operating Generators (by day).			
Saturday	1.55	1,56	2.03
Sunday	1.89	1.75	1.86
Weekday Ave.	2.80	2.64	2.55
Panel B: Costs ($\times 10^{-3}$).			
Maint. B	29	20.2	21.95
Maint. S	20.2	11.8	11.9
Var. B	211.1	235.1	240.4
Var. S	174.8	228.1	215.9
Total/Quantity	0.389	0.452	0.444
Panel C: Quantities and Prices.			
Ave. Q Wkend	93.5	92.0	98.6
Ave. P Wkend	303	325	260
Ave. Q Wkday	185.7	181.8	181.2
Ave. P Wkday	374	401	411

Base Case vs Excess Capacity: AsI & FI

• Maintenance and Shutown.
Base case: the FI equlibrium generates less shutdown and more maintenance.
"Excess" Capacity (more capacity relative to demand) the AsI equilibrium generates less shutdown and more maintenance.

• Weekday vs Weekend.
Base case: AsI vs FI strategies: weekends the AsI equilibrium shuts down more genrators. This enables the firms to signal that their generators will be bid in on the high-priced weekdays.
Excess Capacity: Now the ASI firm no longer distinguishes much between weekend and weekday.

• Prices.
With excess capacity the difference between weekday and weekend prices drops dramatically (to 5.4% in the AsI and 1% in the FI equilibrium) and AsI operation increases on weekend.

- Costs.

Increasing capacity relative to demand the average cost is over 30% lower. Raises questions of what are the capital costs and incentives for private generator construction?

- *Total Surplus:*

Increase in capacity/demand ratio generates a large increase in consumer surplus, and a somewhat smaller total surplus increases. Does increased surplus cover social cost of generator construction? And if so how do we induce the investment?

	Base Case		Excess Capacity	
	AsI	FI	AsI	FI
Panel A: Strategies.				
Firm B: Shutdown and Maintenance.				
Shutdown %	19.96	12.31	41.97	43.75
Maintenance %	10.1	10.9	6.47	6.25
Firm S: Shutdwon and Maintenance.				
Shutdown %	21.48	20.74	53.1	56.4
Maintenance %	9.83	9.91	5.22	4.84
Firm B: Operating Generators (by day).				
Saturday	1.08	1.72	1.03	1.0
Sunday	1.21	1.65	1.03	1.0
Weekday Ave.	1.78	1.78	1.03	1.0
Firm S: Operating Generators (by day).				
Saturday	1,56	2.03	1.21	0.48
Sunday	1.75	1.86	1.20	0.44
Weekday Ave.	2.64	2.55	1.25	1.44
Panel B: Quantities and Prices.				
Ave. Q Wkend	92.0	98.6	33.6	33.1
Ave. P Wkend	325	260	168	175.6
Ave. Q Wkday	181.8	181.2	42.50	42.43
Ave. P Wkday	401	411	177	177

Costs, Consumer Surplus and Total Surplus ($\times 10^{-3}$) .

	Base Case		Excess Capacity	
	AsI	FI	AsI	FI
Average Cost	.452	.444	.290	.282
CS*	581.5	595.0	1,316	1,311
Total Surplus	288.9	301.4	1,374	1,373

* CS= these numbers plus 58,000.

** Total Surplus = these numbers plus 59,000.

Conclusions.

• There is a need for increased research on the dynamics of market outcomes.

• The framework used for this analysis ought probably to require less of both the agent and the analyst then does "Bayesian Perfect" notions of equilibria.

• Ultimately, that framework will have to integrate the analysis of the reactions to changes in institutions with an analysis of policies for states that are observed repeatedly.

• Some empirical knowledge of how firms react to changes in their environment would be extremely helpful both in; (i) formulating an appropriate model, and (ii) in selecting equilibria when analyzing counterfactuals.

• Policies for situations that are observed repeatedly should be modelled to satisfy the conditions of an "Experience Based Equilibrium", and if more stringent equilibrium conditions are justified they should be imposed as they will result in a more precise analysis.

• Most changes we are likely to want to analyze are not as dramatic as setting up a whole new market, and adaptation processes used intensively in our neighboring OR and CS fields, processes that do not explicitly require explicit experimentation, might generate good approximations to them.

• Our editors and referees should remember that the world is too complicated, and our data and knowledge of appropriate modeling techniques too limited, to expect us to analyze models that satisfy the rigor of formal statistical testing procedures. The more useful test is a test of relevance: does the analysis generate a better approximation than the alternatives currently available?