

Methodological Issues in Analyzing Market Dynamics.

**Cowles Lecture
Econometric Society.
Summer Meetings,
Minneapolis, 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 the likely causes of historical events, or 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.

Claim 1. The tools developed for the analysis of market allocations conditional on the state variables of the problem (characteristics of products marketed, cost functions, ...) pass a market test for success as:

- (i) They have been incorporated into applied work in virtually all of economics that deals with market allocations (especially where demand and/or productivity analysis is needed),
- (ii) They are used by public agencies, consultancies and to some extent by firms and
- (iii) They fit surprisingly well (at least in the cross-section).

Note. There are improvements still being done, and important new work in analyzing equilibrium allocations in situations where Nash in prices or quantities seems less relevant; e.g. vertical markets characterized by multi-agent bargaining with externalities, platform markets,...

E.g. of Fit: Pricing Behavior. Wollman's dissertation (commercial trucks). Estimate BLP demand, regress Nash markup on instruments to get $\widehat{markup}$ ($R^2=.44$ or $.46$ with time dummies; sophisticated IV would do better). Look to fit & whether 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.

Note. Of the 14% of variance unexplained, over half (8%) picked up by time dummies. There are level shifts across periods.

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.

The initial frameworks by our theory colleagues made assumptions which insured that the

- state variables evolve as a Markov process
- and the equilibrium is some form of Markov Perfection (no agent has an incentive to deviate at any value of the state variables).

E.g. Maskin and Tirole (1988). Except in situations involving active experimentation and learning (where policies are transient), applied work is likely to stick with the assumption that market dynamics are generated by a time homogenous Markov process (of finite order).

Theory may lead you to worry about the (finite order) Markovian assumption (dynamic games of asymmetric 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.

The type of rationality built into Markov Perfection is more questionable; though useful in the simple models used by our theory colleagues to explore possible outcomes in a structured way.

Empirical work on dynamics proceeded in a similar way to the way we did in static analysis; we took the Markov Perfect framework and tried to incorporate the institutions that seemed necessary to analyze actual markets.

What was the result? The dynamic framework we inherited became unwieldy when confronted with the complexity of the institutions we were trying to model. The complexity of the dynamic framework 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.

As we shall see these are useful in a broader context, but they were simply not enough to allow us to convert Markov Perfection into a sufficiently realistic framework for the empirical analysis of market dynamics. As a result many dynamic issues are analyzed in a much less rigorous way than static issues (e.g; empirical work on the dynamics of mergers).

Claim 2. When we turn to empirical work (in contrast to applied theory) our dynamic tools have not passed beyond the hands of a few diligent I.O. researchers. This severely

- limits our ability to do dynamic analysis of market outcomes, and
- leads to a question of whether some other notion of equilibria will better approximate agents' behavior.

I want to **suggest** that 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.

I.e. If we abandon Markov Perfection can we both

- better approximate agents' behavior and,
- enlarge the set of dynamic questions we are able to answer.

The complexity issue. When we try to incorporate what seems to be essential institutional background we find

- 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)

How demanding is this? 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 a doubly nested fixed point, or learn and retain, policies from each distinct information set.

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 asymmetric information: the “perfectness” notion would then lead us to a “Bayesian” Markov Perfect solution.

Problem. The burden of computing its strategies insures that they will not be directly computed by either agents or the analyst 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.


Of course agents might obtain these policies, or at least policies which maintain some of the logical features of Bayesian Perfect policies, in other ways. An obvious possibility is that they are learned from combining data on past behavior with market outcomes. They would have to learn about;

- primitives (some empirical work on this),
- the likely behavior of their competitors, and
- 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.

An ongoing study by U. Doraszelski, G. Lewis and myself of bids from the date the British Electric Utility market for frequency response opened, addresses this question. The conclusions we are reasonably confident on to date are (see the figures)

- The bids do eventually converge, and they converge to what looks like a Nash equilibrium (in 2009), and
- In its initial stages, the learning process is complex, involves experimentation, and differs among firms.
- The smaller changes that occurred in the first half of 2010 (Drax signs long-term contract with NG) seems more structured and proceeds to an equilibrium much quicker.



— Aberthaw

— Drax

— Rats


— Cottam

— Eggborough

— Seabank

— Connah's Quay

— Peterhead



Rest of Talk.

- Unfortunately, I have little to say about periods of “active” learning and experimentation. However I do have a suggestion for modelling more stable environments.
- For these environments, I will introduce
 - (i) a notion of equilibrium that is less demanding than Markov Perfect for both the agents, and the analyst, and show how to
 - (ii) compute the equilibrium and
 - (iii) estimate off of equilibrium conditions.
- Consider what can be done about multiple equilibria.
- Provide a computed example of this equilibrium (context of electric utility generation).

I start with strategies that are “rest points” to a dynamical system. Later I will consider institutional change, but only changes where it is reasonable to model responses to the change with a simple reinforcement learning process (a process which does not require assigning value to experimentation). This makes my job easier.

- Easier because the strategies at the rest point likely satisfy a Nash condition of some sort; else someone has an incentive to deviate.
- However it still opens up at least one important question – what is the form of the Nash Condition?

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

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 at least consistent with what they observe.

Note. In situations where it might be reasonable to assume more than this, as we will detail below, there is good reason to impose any other restrictions available.

Notation & the Markov Assumption.

- Denote the information set of firm i in period t by $J_{i,t}$. $J_{i,t}$ will contain both public (ξ_t) and private ($\omega_{i,t}$) information, so $J_{i,t} = \{\xi_t, \omega_{i,t}\}$.
- 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.
- Policies, say $m_{i,t} \in \mathcal{M}$, will be functions of $J_{i,t}$. For simplicity we will assume $\#\mathcal{M}$ is finite, and that it is a capital accumulation game, i.e. $\forall (m_i, m_{-i}) \in \mathcal{M}^n$, & $\forall \omega \in \Omega$

$$P_\omega(\cdot | m_i, m_{-i}, \omega) = P_\omega(\cdot | m_i, \omega).$$

ξ = the information that each agent has on its competitors' ω 's and common demand and cost conditions which evolve as an exogenous Markov process.

- So a “state” of the system, is

$$s_t = \{J_{1,t}, \dots, J_{n_t,t}\} \in \mathcal{S},$$

$\#\mathcal{S}$ is finite, and any set of policies will insure that s_t will wander into a recurrent subset of $\mathcal{S}$, say $\mathcal{R} \subset \mathcal{S}$, in finite time, and after that $s_{t+\tau} \in \mathcal{R}$ w.p.1 forever.

- 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

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

- and of expected profits be

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

If $m^*(J_i)$ is the policy chosen at J_i , then **our assumptions imply**

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

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

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

where, if $p^e(\cdot)$ provides the empirical probability
(the number of times the event occurs)

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

and

$$\left\{ p^e(J_{-i}|J_i) \equiv \frac{p^e(J_{-i}, J_i)}{p^e(J_i)} \right\}_{J_{-i}, J_i},$$

while

$$\left\{ p^e(J'_i|J_i) \equiv \frac{p^e(J'_i, J_i)}{p^e(J_i)} \right\}_{J'_i, J_i} \cdot \spadesuit$$

“Restricted Experience Based Equilibrium”

These are the conditions of a REBE (Fershtman and Pakes, (2012); for related earlier work see Fudenberg and Levine (1993)). Bayesian Perfect satisfy them, but so do weaker notions. We now turn to its :

- (i) computational and estimation properties,
- (ii) overcoming multiplicity issues,
- (iii) and then to an example.

Computational Algorithm. Iterative “reinforcement learning” algorithm.

- Might be viewed as an approximation to a learning process which makes it a candidate to analyze perturbations to the environment as well as to compute equilibrium.
- Does not (necessarily) generate a curse of dimensionality; makes it relatively easy for an analyst as well as an agent to use.

Iterations defined by

- A location, say $L^k = (J_1^k, \dots, J_{n(k)}^k) \in \mathcal{S}$: is the information sets of agents active.
- Objects in memory (i.e. M^k):
 - (i) perceived evaluations, W^k ,
 - (ii) No. of visits to each point, h^k .

So algorithm must update (L^k, W^k, h^k) . Computational burden determined by; memory constraint, and compute time. The memory burden depends on how we structure the $W(\cdot|\cdot)$. I will present the form which is simplest notationally, but it is not necessarily the most efficient way.

Update Location.

- Calculate policies for all agents active at L^k to maximize the agents' values conditional on the evaluations in memory (W^k). I.e.

$$\max_{m \in \mathcal{M}} W^k(m|J_i)$$

- Take random draws on outcomes conditional on those policies and L^k (in our example on η conditional on m) and use them to form $\{J_i^{k+1}\}_i$.

Update W^k .

Assuming $b(m_{-i})$ is publically observed and the $W^k(\cdot)$ are the agent's evaluations, the agent's ex poste perception of what its value would have been had it chosen m would be

$$V^{k+1}(J_i^k, m) = \pi(\omega_i^k, m, b(m_{-i}^k), d^k) + \max_{\tilde{m} \in M} \beta W^k(\tilde{m} | J_i^{k+1}(m)),$$

where $J_i^{k+1}(m)$ is what the $k + 1$ information would have been given m and *competitors actual play*.

Treat $V^{k+1}(J_i^k)$ as a random draw from the possible realizations of $W(m|J_i^k)$, and update W^k as we update averages

$$W^{k+1}(m|J_i^k) - W^k(m|J_i^k) = \frac{1}{h^k(J_i^k)} [V^{k+1}(J_i^k, m) - W^k(m|J_i^k)].$$

or

$$W^{k+1}(m|J_i^k) = \frac{1}{h^k(J_i^k)} V^{k+1}(J_i^k, m) + \frac{(h^k(J_i^k) - 1)}{h^k(J_i^k)} W^k(m|J_i^k)$$

Other weights might be more efficient. Note if we are in equilibrium we tend to stay there, i.e. if $*$ indicates equilibrium values and policies, then $\forall J_i \subset s \in \mathcal{R}$

$$E[V^*(J_i^k, m^*)|W^*] = W^*(m^*|J_i^k).$$

Notes.

- Agents (as well as the analyst) could (at least in principal) use the algorithm to find equilibrium policies or adjust to perturbations in the environment.
- The algorithm has no curse of dimensionality in either calculating continuation values or in the cardinality of the points updated.
 - (i) Computing continuation values is replaced by averaging two numbers.
 - (ii) The algorithm eventually wanders into and stays in $\mathcal{R}$ whose cardinality need not increase with the dimension of the state space (though it often goes up linearly in that dimension).
- The algorithm uses the computational advances of the stochastic approximation literature and as in that literature it can be augmented to use functional form approximations where needed (“TD learning”).

Computational Properties.

- Testing. The algorithm does not necessarily converge, but a test for convergence exists and does not involve a curse of dimensionality.
- The test is based on simulation. It produces a consistent estimate of an $L^2(P(\mathcal{R}))$ norm of the percentage bias in the implied estimates of $V(m, J_i)$; where $P(\mathcal{R})$ is the invariant measure on the recurrent class.
- I am going to assume that there is a point which all points communicate with. This insures that the recurrent class associated with a given estimate of $\mathcal{W} \equiv \{W(m, J_i)\}_{m, J_i \in s \in \mathcal{S}}$ is unique. One can often prove the existence of such a point theoretically, and one can always examine whether the condition holds for a computed $\mathcal{W}$.

Estimation.

The equilibrium 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 different 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 equilibrium 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.
- The conditional (on J_i) expectation of the difference in discounted profits between the simulated and actual path from the period of the deviation to the random stopping time, should, when evaluated at the true parameter vector, be positive. This enables the development of moment inequalities for estimation (see Pakes, 2010).

Multiplicity.

- $\mathcal{R}$ contains both “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.
- Our conditions only insure that perceptions of outcomes are consistent with the results from actual play at interior points. Perceptions of outcomes for some feasible (but inoptimal) policy at boundary points are not tied down by actual outcomes.
- MPBE are a special case of REBE and they have multiplicity. Here differing perceptions at boundary points can support a (possibly much) wider range of equilibria.

Narrowing the Set of Equilibria.

- In any empirical application 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.
- Allow firms to experiment with $m_i \neq m_i^*$ at boundary points. Leads to a stronger notion of, and test for, equilibrium (insure all perceptions are consistent with the results from actual play at boundary points).
- Use external information to constrain perceptions of the value of outcomes outside of $\mathcal{R}$.

Boundary Consistency.

Let $B(J_i|\mathcal{W})$ be the set of actions at $J_i \in s \in \mathcal{R}$ which could generate outcomes which are not in the recurrent class. If $\text{supp}[p(\cdot|m, \omega_i)]$ is the support of $p(\cdot|m, J_i)$ and ϕ is the null set then

$$B(J_i|\mathcal{W}) \equiv$$

$$\left\{ (m \in \mathcal{M}, J_i \in \mathcal{R}) : \text{supp}[p(\cdot|J_i, m)] \cap J'_i \in s \notin \mathcal{R} \neq \phi \right\}.$$

Let $B(\mathcal{W}) = \cup_{J_i \in \mathcal{R}} B(J_i|\mathcal{W})$ and τ index future periods. Then the extra condition needed to insure “Boundary Consistency” is:

Extra Condition. $\forall (m, J_i) \in B(\mathcal{W})$

$$E \left[\sum_{\tau=0}^{\infty} \delta^\tau \pi(m(J_{i,\tau}), m(J_{-i,\tau})) | J_i = J_{i,0}, \mathcal{W} \right] \leq W(m^*|J_i),$$

where $E[\cdot|J_i, \mathcal{W}]$ takes expectations over future states starting at J_i and using the policies generated by $\mathcal{W}$. ♠

Testing for Boundary Consistency

Assuming unique $\mathcal{R}$. From each $(m, J_i) \in B(\mathcal{W})$ simulate independent sample paths. Indexing the periods of a path by τ , terminate it the first time $J_i \in s \in \mathcal{R}$, say τ^* . The path's estimate of $W(m|J_i)$ is

$$\tilde{W}(m|J_i) \equiv \sum_{\tau=0}^{\tau^*} \delta^\tau \pi(m(J_{i,\tau}), m(J_{-i,\tau})) + \delta^{\tau^*} W(m^*|J_{i,\tau^*}),$$

with mean and variance; $\bar{W}(m|J_i)$, $V[\bar{W}(m|J_i)]$.

Let $f(x)_+ = \max[0, f(x)]$, and

$$T(m|J_i) \equiv \frac{[\bar{W}(m|J_i) - W(m^*|J_i)]_+}{W(m^*|J_i)}$$

Now use a one-sided test of

$$H_0 : \frac{\sum_{(m,J_i) \in B(\mathcal{W})} T(m|J_i)}{\sqrt{\sum_{(m,J_i) \in B(\mathcal{W})} V[T(m|J_i)]}} = 0,$$

where $V[T(m|J_i)]$ is the variance of $T(m|J_i)$.

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 electricity 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.

- $b(m_i) : m_i \rightarrow \{0, b_i\}^{n_i}$ 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

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

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

$$m_{i,r,t} = 0 \Rightarrow \omega_{i,r,t+1} = \omega_{i,r,t},$$

$$m_{i,r,t} = 1 \Rightarrow \omega_{i,r,t+1} = \bar{\omega}_{i,r} \text{ (}\bar{\omega}\text{=restart state),}$$

$$m_{i,r,t} = 2 \Rightarrow \omega_{i,r,t+1} = \omega_{i,r,t} - \eta_{i,r,t}$$

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. $s_{i,t} = (J_{1,t}, \dots, J_{n_t,t}) \in \mathcal{S}$,
and

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

where

- $\omega_{i,t}$ represents private information
- and ξ_t is public information (shared by all).
Example $\xi_t = \{b(m_{1,\tau}), b(m_{2,\tau}), d_\tau\}_{\tau \leq t}$, and
knowledge of $\omega_{-i,t}$ the last period of revelation (happens every T periods).

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 equilibrium 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 generators. 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 than 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. “Adaptation” processes, like reinforcement learning, might be adequate for reactions to changes that do not induce calculated experimentation.
- A start for equilibrium conditions at situations that are observed repeatedly are those of “Experience Based Equilibrium”. If more stringent equilibrium conditions are justified they should be imposed as they will result in a more precise analysis.