

Methodology for Empirical Analysis of Market Dynamics.

CRESSE Conference on
Competition and Regulation.
Rethymnon Crete, July 2015.

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 background 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,
 - (ii) **equilibrium conditions** when we need to solve for variables that firms could change in response to the environmental change of interest.

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

- Each agent's actions affect all agents' payoffs, 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.

The difficulties encountered in incorporating sufficient heterogeneity and/or using equilibrium conditions differed between “static” and “dynamic” models.

I use the textbook distinction here:

- (i) static models solve for profits conditional on state variables, and
- (ii) dynamics analyzes the evolution of those state variables (and through that the evolution of market structure).

Empirical Claim 1. The tools developed for the analysis of market allocations conditional on the “state variables” of the problem (characteristics of products marketed, cost determinants, ...) pass a market test for success. Reasons for claim:

- (i) They have been incorporated into applied work in virtually all of economics that deals with market allocations (especially where productivity and/or demand is needed),
- (ii) They are used by public agencies, consultancies and to some extent by firms and
- (iii) They do surprisingly well, both in fit and in providing a deeper understanding of empirical phenomena, as the example we now go to demonstrates.

Note. There are improvements still being done, and important new work in analyzing equilibrium allocations in markets where Nash in prices or quantities seems inappropriate; e.g. vertical markets, 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. Level shifts (time dummies) are 8% of the 14% of unexplained variance.

E.g. Productivity analysis. Olley and Pakes (1996). Equilibrium conditions used to control for simultaneity in input demands and selection due to exit in "new" micro production function data sets. Estimates enable a decomposition of the effect of environmental changes into

- the efficiency of the allocation of output among firms, and
- individual firm productivities.

Initial e.g: productivity impacts of breakup of A.T.&T (telecom equipment). Sharp increase in productivity after the breakup due to a re-allocation of output to more productive firms (largely as a result of a reallocation of capital).

Ongoing Work. De Loecker (2014) and coauthors. Use production function estimation (above) and input equilibrium conditions to separate changes in revenue into:

- changes in markups and
- changes in costs.

Early e.g: extent of pass-through after lowering of Indian tariffs. Most of the reduction in input costs went to higher markups.

” Static” Product Repositioning.

Realistic analysis of product repositioning in industries where a firms’ product characteristics can be changed as quickly as price. Recent e.g.’s: semiconductors (Eisenberg, 2014, Nosko, forthcoming), ” modular” trucks (Wollman 2015). These papers analyze the change in products marketed in response to; innovation, counterfactual mergers and/or counterfactual exit.

Tools required: (i) demand analysis (as above), (ii) inequality estimators to estimate fixed costs of repositioning (Ciliberto & Tamer, 2012; Pakes Porter Ho and Ishii, 2015).

Inequality estimators. If we

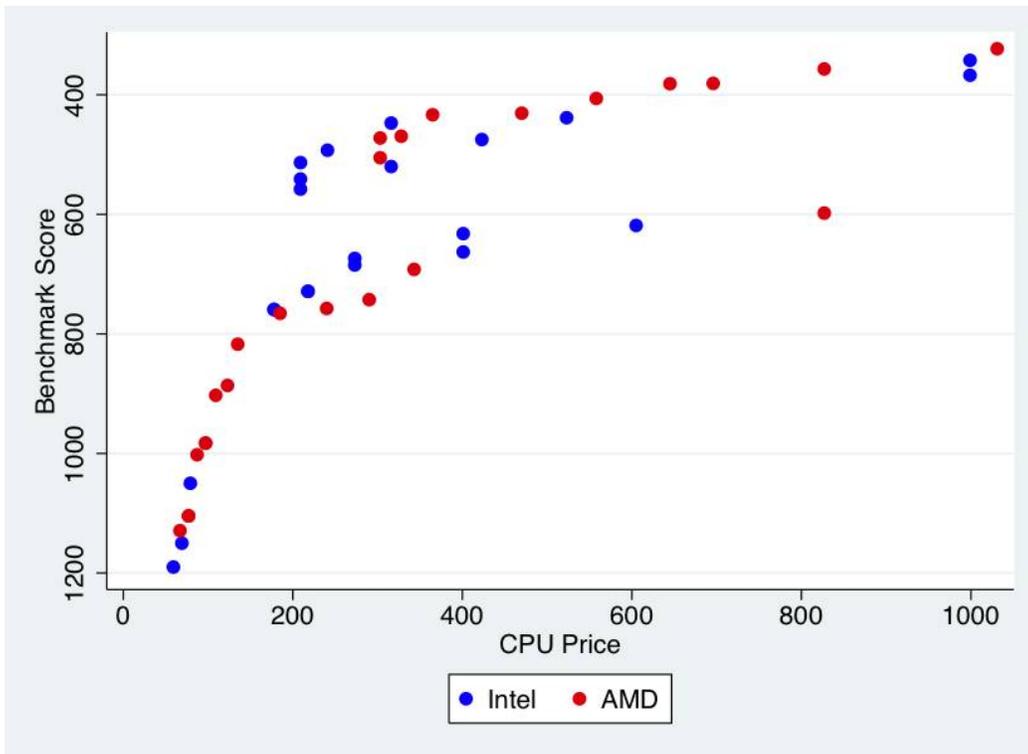
(i) Reposition: increment in expected profit must be greater than the FC of repositioning,

(ii) Do not reposition: the increment in expected profit less than the FC of repositioning.

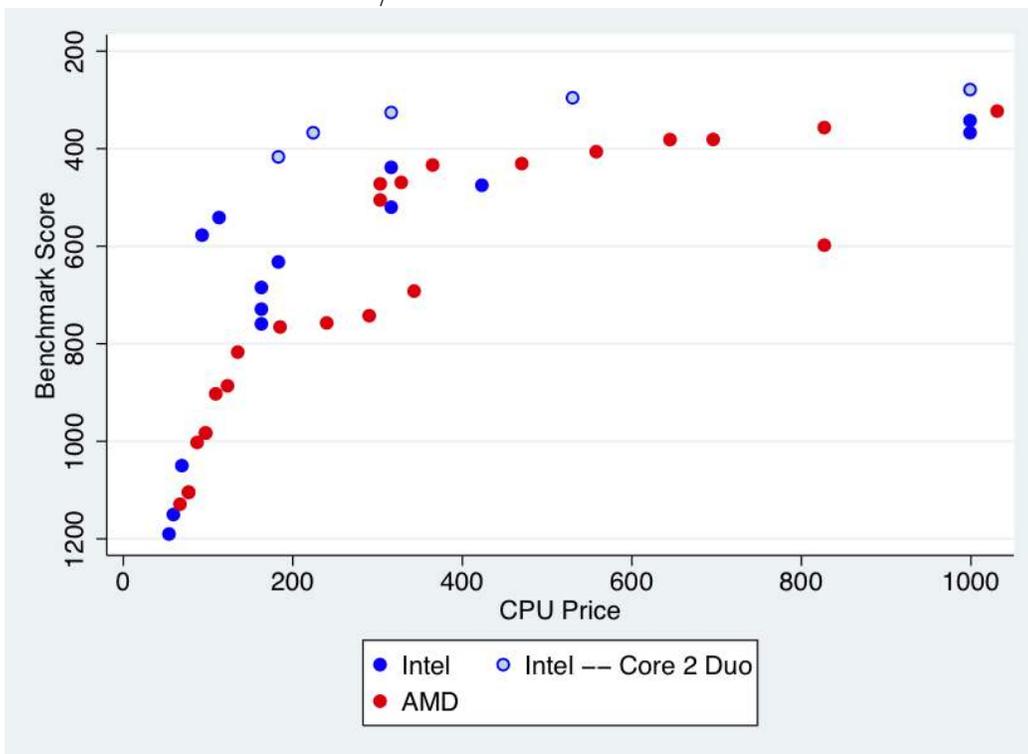
E.g.: Nosko, Harvard thesis. Studies competition in pc chips (AMD/Intel). Product characteristics change more frequently than price of existing products. Analyzes: (i) the introduction of Intel's Core 2 Duo (July 2006; it lowers the cost of producing all chips, and increases max performance), and (ii) considers a counterfactual AMD/Intel merger.

Graphs: Benchmarks refer to time required to complete tasks. Price=published list prices.

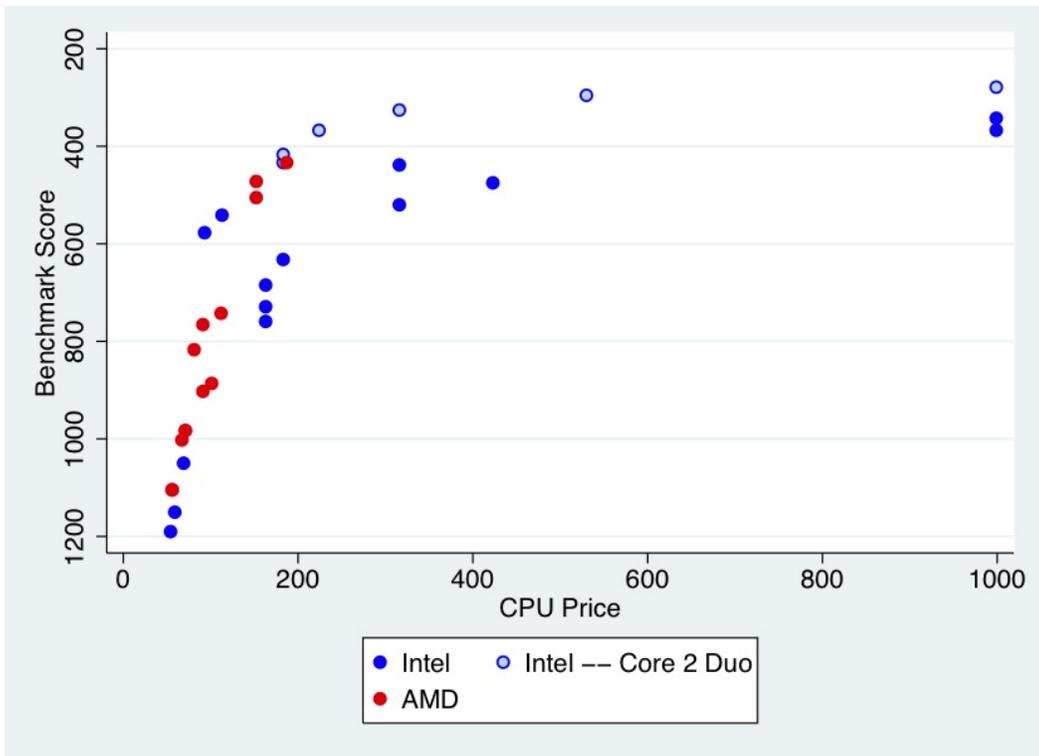
- In June: Intense competition at high end high-end between Intel and AMD.
- October 2006. Intel had begun introducing the Core 2 duo chips. AMD immediately



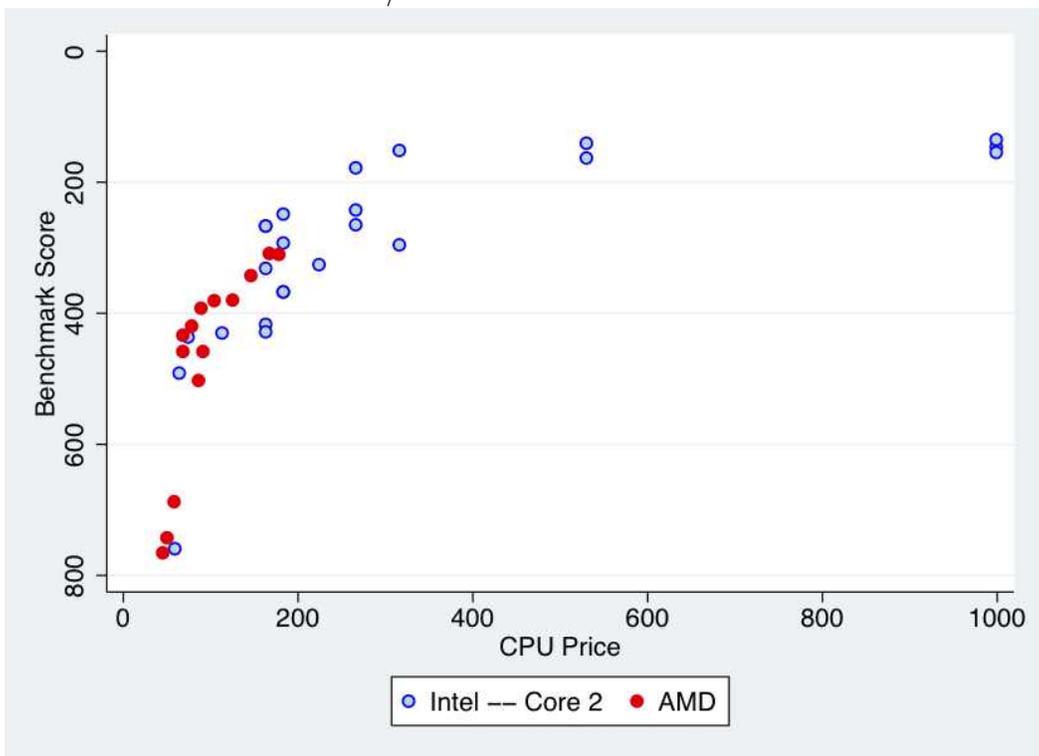
Price/Performance – June 2006



Price/Performance – July 2006



Price/Performance – Oct 2006



Price/Performance – January 2008

withdraws all its chips from the high end of the market, and focuses on competing on the low end.

- January 2008. Within a year and half all of Intel's chips were Core 2 duo chips. The benchmark performance scores at the high end had gone up, and the only competition with AMD is at the low end.

Profit from Core 2 duo comes from intel dominating high end and pushing up markups their. Merger counterfactual: intel withdraws virtually all of its mid range chips and pushes up markup on high end chips.

Implication. The returns to either innovation or merger could not be analyzed without product repositioning.

What about market dynamics?

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

1. state variables evolve as a Markov process
2. 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) and Ericson and Pakes (1995). We now consider each of these in turn.

On the Markov Assumption. Except in situations involving active experimentation and learning (where policies are transient), applied work is likely to stick with the assumption that states evolve as a time homogenous Markov process of finite order. Reasons:

- It is a convenient and fits the data well.
- Realism suggests information access and retention conditions limit the memory used.
- We can bound unilateral deviations (similar to Weintraub, 2014), and have conditions which insure those deviations can be made arbitrarily small by letting the length of the kept history grow (White and Scherer, 1994).

On 2: Perfection. The type of rationality built into Markov Perfection is more questionable; even though it has been useful in the simple models used by our theory colleagues to explore possible outcomes in a structured way. We come back to this below.

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

The Result. Though the MP framework was useful in controlling or dynamic issues in the analysis of several problems (e.g. controlling for selection in analyzing productivity) it became unwieldy when confronted with the task of analyzing market dynamics. This because of the complexity of the institutions we were trying to model. The difficulties 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; e.g.'s Bajari et. al (BBL), Pakes et. al. (POB).
- the use of approximations and functional forms for primitives which enabled us to compute equilibria quicker and/or with less memory requirements than in the original Pakes and McGuire (1994) article (Doraszelski and Judd, 2011; Pakes and McGuire, 2001).

The underlying ideas:

- are useful under other equilibrium assumptions (see below), and
- enabled an expansion of computational dynamic theory and a deeper understanding of a number of dynamic issues.
 - (i) LBD; Besanko, Doraszelski, Satterthwaite 2010.
 - (ii) Market responses to merger policy; Mermelstein, Satterthwaite, Whinston 2014.
 - (iii) Collusion and consumer welfare; Fershtman and Pakes, 2000.

However empirical work on dynamic issues remains extremely limited and what is done, though perhaps better than other analysis of the issues at hand, typically requires use of assumptions that are distinctly more questionable than those used in static work.

I want to start from the premise that the complexity of Markov Perfection not only limits our ability to do dynamic analysis of market outcomes it also

- leads to a question of whether some other notion of equilibria will better approximate agents' behavior.

So the fact that 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 equilibrium will better approximate agents' behavior.

Question. If we abandon Markov Perfection can we both

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

The complexity issue. When we try to incorporate "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 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.

Is "Bayesian MP" more realistic?

- It decreases the information access and retention conditions making it more realistic.
- However it increases the burden of computing the policies significantly over the burden of computing in symmetric information MPE models. 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.

Could agents learn these policies? Or at least policies which maintain some of the logical features of Bayesian Perfect policies, 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 (monthly) bids from the date the British Electric Utility market for frequency response opened (late 2005), 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 early 2009).
- In its initial stages (until the end of 2006), the learning process is complex, involves experimentation, and differs among firms.

- There is a second stage (in 2007 & 2008) which exhibits fairly smooth transitions to what looks like a stable equilibrium.
- After 2009 there were a series of smaller changes in the environment (that were largely absent in the earlier periods) involving forward contracts, entry, and fuel prices. Perhaps most surprisingly, these caused just minor disruptions in bids placed and almost immediate reversion to what looked like an equilibrium.

Rest of Talk.

- Unfortunately, I have little to say on “active” experimentation periods; on modeling beliefs on the value of different experiments.
- For more stable environments I
 - (i) incorporate asymmetric information (for both realism and theoretical consistency),
 - (ii) but introduce a notion of equilibrium that is easier than Bayesian Perfect for both the agents and the analyst to use, and
 - (iii) then show how to compute the equilibrium and estimate off of equilibrium conditions.
- Consider restrictions that mitigate multiplicity issues.

- Provide a computed example of this equilibrium (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 (I do not consider changes that lead to active experimentation). This makes my job much easier because:

- Strategies at the rest point likely satisfy a Nash condition of some sort; else someone has an incentive to deviate.
- However it still leaves opens the 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 I believe 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. It might be reasonable to assume more than this: that agents (i) know and/or (ii) explore, properties of outcomes of states not visited repeatedly. I come back to this below.

Formalization of Assumptions.

- 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 can be adequately approximated by one); and only a finite number of firms are ever simultaneously active.
- Policies, say $m_{i,t} \in \mathcal{M}$, will be functions of $J_{i,t}$. For simplicity assume $\#\mathcal{M}$ is finite, and that it is a simple capital accumulation game, i.e. $\forall (m_i, m_{-i}) \in \mathcal{M}^n$, & $\forall \omega \in \Omega$

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

So the public information, ξ , is used to predict competitor behavior and common demand and cost conditions (these evolve as an exogenous Markov process).

- A “state” of the system, is

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

$\#\mathcal{S}$ is finite. \Rightarrow 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.

- The agents does not know s_t , or calculate policies for its components. Symmetry would reduce the dimension of the state space to $\#\mathcal{J}$, otherwise $N \times \#\mathcal{J}$.

- 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

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

Our assumptions imply:

- Each agent choses an action which maximizes its perception of its expected discounted value, and
- For those states that are visited repeatedly (are in \mathcal{R}) these perceptions are consistent with observed outcomes.

Formally

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

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

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

where, if $p^e(\cdot)$ provides the empirical probability (the fraction of periods 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

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

“Experience Based Equilibrium”

These are the conditions of a (restricted) EBE (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. Asynchronous “Reinforcement learning” algorithm (Pakes and McGuire, 2001). Two advantages:

- (i) No “curse of dimensionality” (should be able to increase the number of state variables),
- (ii) Mimics a learning process (makes it a candidate to analyze perturbations to the environment, as well as to compute equilibrium).

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 .

Must update (L^k, W^k, h^k) . Computational burden determined by; memory constraint, and compute time. I use a simple (not necessarily) optimal structure to memory.

Update Location.

- Calculate “greedy” policies for each agent

$$m_{i,k}^* = \arg \max_{m \in \mathcal{M}} W^k(m | J_{i,k})$$

- Take random draws on outcomes conditional on $m_{i,k}^*$: i.e. if we invest in “payoff relevant”

$\omega_{i,k} \in J_{i,k}$, draw $\omega_{i,k+1}$ conditional on $(\omega_{i,k}, m_{i,k}^*)$.

- Use outcomes to update $L^k \rightarrow L^{k+1}$.

Update W^k .

- “Learning” interpretation: Assume agent observes $b(m_{-i})$ and knows the primitives;

$\pi_i(\cdot), p(\omega_{i,t+1} | \omega_{i,t}, m_{i,t})$.

- Its ex poste perception of what its value would have been had it chosen m is

$$V^{k+1}(J_{i,k}, m) =$$

$$\pi(\omega_{i,k}, m, b(m_{-i,k}), d_k) + \max_{\tilde{m} \in M} \beta W^k(\tilde{m} | J_{i,k+1}(m)),$$

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 in stochastic integration (Robbins and Monroe, 1956)

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

or

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

(other weights are more efficient, it would be good to know how to aggregate states)

Notes.

- If we have equilibrium valuations we tend to stay their, i.e. if * designates equilibrium

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

- To learn equilibrium values we need to visit points repeatedly; only likely for states in \mathcal{R} .
- Agents (not only the analyst) could use the algorithm to find equilibrium policies or adjust to perturbations in the environment.
- Algorithm has no curse of dimensionality.
 - (i) Computing continuation values: integration is replaced by averaging two numbers.
 - (ii) States: algorithm eventually wanders into \mathcal{R} and stays there, and $\#\mathcal{R} \leq \#\mathcal{J}$.
- The stochastic approximation literature often augments this with functional form approximations (“TD learning”; Sutton and Barto, 1998). There is a question of how to reduce the dimension of the state space. Could either try to (i) mimic computational results; or (ii) provide accurate predictions of how firms behave.

Convergence and Testing.

- Testing. The algorithm does not necessarily converge, but a test for convergence exists and does not involve a curse of dimensionality (Fershtman and Pakes, 2012).
- 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.
- **Basis.** Simulate sample path. Eventually will visit each $J_i \in \mathcal{R}$ many times. Store realized mean value ($\tilde{W}(m|J_i)$) and variance ($\tilde{V}(\tilde{W}(m|J_i))$) of playing m at J_i . $(\tilde{W}(m|J_i) - W(m|J_i))^2$ is the MSE of $\tilde{W}(m|J_i)$ as an estimate of $W(m|J_i)$.

$$\%Bias(m|J_i)^2 = \frac{(\tilde{W}(m|J_i) - W(m|J_i))^2}{W(m|J_i)^2} - \frac{\tilde{V}(\tilde{W}(m|J_i))}{W(m|J_i)^2}.$$

Computational Issues.

- **Excessive visits** to points that have a high probability in the invariant measure. We update certain points many more times than is necessary. Test is point-wise: put "converged" points in a separate list and do not update them thereafter (though include them in subsequent tests).
- **Aggregation over states.** For realistic problems there are more states than are necessary for the accuracy of computed equilibrium (or that agents are likely to use). Question: how do we aggregate states? For empirical problems we might start by finding adequate approximations to the determinants of the dynamic controls.

Details. Any fixed W , say \tilde{W} , generates policies which define a finite state Markov process for $\{s_t\}$. Gather the transition probabilities into the Markov matrix, $Q(s', s|\tilde{W})$.

To test if the process satisfies our equilibrium conditions need:

- (i) a candidate for \mathcal{R} , and checks for
- (ii) optimality of policies and
- (iii) consistency of W .

Candidate for $\mathcal{R}(\tilde{W})$. Start at any s^0 and use $Q(\cdot, \cdot|\tilde{W})$ to simulate a sample path $\{s^j\}_{j=1}^{J_1+J_2}$. Let $\mathcal{R}(J_1, J_2, \cdot)$ be the set of states visited at least once between $j = J_1$ and $j = J_2$.

$$(J_1, J_2) \rightarrow (\infty, \infty), \quad \& \quad J_2 - J_1 \rightarrow \infty$$

$$\Rightarrow \mathcal{R}(J_1, J_2, \cdot) \rightarrow \tilde{\mathcal{R}},$$

a recurrent class of $Q(\cdot, \cdot|\tilde{W})$ (C1 satisfied).

C2 (optimality of policies). Satisfied by construction, since we use the policies generated by \tilde{W} to form $Q(\cdot, \cdot|\tilde{W})$.

C3 (consistency of \tilde{W} with outcomes). Does

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

$(\forall J_i \in s \in \mathcal{R}.)?$

Direct summation. Computationally burdensome; indeed brings the curse of dimensionality back in.

Alternative. Check for consistency of simulated sample paths with evaluations.

- Start at $s_0 \in \mathcal{R}$ and forward simulate. At each J_i compute perceived values (our $V^{k+1}(\cdot)$), keep track of the average and the sample variance of those simulated perceived values, say

$$\left(\hat{\mu}(\tilde{W}(m^*(J_i)|J_i)), \hat{\sigma}^2(\tilde{W}(m^*(J_i)|J_i)) \right).$$

- Let $E(\cdot)$ take expectations over the simulated random draws (where draws will be indexed by a tilde), let l index locations, and note that we can compute \mathcal{T}_l , where

$$\begin{aligned}\mathcal{T}_l &\equiv E\left(\frac{\hat{\mu}(\tilde{W}_l) - \tilde{W}_l}{\tilde{W}_l}\right)^2 \\ &= E\left(\frac{\hat{\mu}(\tilde{W}_l) - E[\hat{\mu}(\tilde{W}_l)]}{\tilde{W}_l}\right)^2 + \left(\frac{E[\hat{\mu}(\tilde{W}_l)] - \tilde{W}_l}{\tilde{W}_l}\right)^2. \\ &= \%Var(\hat{\mu}(\tilde{W}_l)) + \%Bias^2(\hat{\mu}(\tilde{W}_l)).\end{aligned}$$

- \mathcal{T}_l is observed, as is f_l , the fraction of visits to l . As the number of simulation draws grows

$$\sum_l f_l \left(\frac{\hat{\sigma}^2(\tilde{W}_l)}{\tilde{W}_l^2}\right) - \sum_l f_l \left(\frac{\hat{\mu}(\tilde{W}_l) - E[\hat{\mu}(\tilde{W}_l)]}{\tilde{W}_l}\right)^2 \rightarrow_{a.s.} 0,$$

\Rightarrow

$$\sum_l f_l \mathcal{T}_l - \sum_l f_l \left(\frac{\hat{\sigma}^2(\tilde{W}_l)}{\tilde{W}_l^2} \right) \rightarrow_{a.s.} \sum_l f_l \left(\frac{E[\hat{\mu}(\tilde{W}_l)] - \tilde{W}_l}{\tilde{W}_l} \right)^2,$$

an $L^2(\mathcal{P}_{\mathcal{R}})$ norm in the percentage bias ($\mathcal{P}_{\mathcal{R}}$ is the invariant measure associated with (\mathcal{R}, \tilde{W})).

Estimation.

- Need a candidate for J_i . Either:
 - (i) empirically investigate determinants of controls (determinants of controls), and/or
 - (ii) ask actual participants.
- Does not require nested fixed point algorithm. Use estimation advances designed for MP equilibria (POB or BBL), or a perturbation (or “Euler” like) condition (below).

Euler-Like Condition.

- With asymmetric information the equilibrium condition

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

is an inequality which can generate (set) estimators of parameters.

- J_i 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 .
- If $m^*(J^2)$ results in outcomes in \mathcal{R} , we can simulate a sample path from J^2 using only observed equilibrium play. The Markov property insures 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 yields moment inequalities for estimation as in Pakes, Porter, Ho and Ishii (2015), 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 (restricted) EBE 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.
- Use external information to constrain perceptions of the value of outcomes outside of \mathcal{R} . If available use it.
- Allow firms to experiment with $m_i \neq m_i^*$ at boundary points (as in Asker, Fershtman, Jihye, and Pakes, 2014). Leads to a stronger notion of, and test for, equilibrium. We insure that perceptions are consistent with the results from **actual play** for each **feasible** action at boundary points (and hence on \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 (so J_i is a boundary point) and $B(\mathcal{W}) = \cup_{J_i \in \mathcal{R}} B(J_i|\mathcal{W})$. Then the extra condition needed to insure “Boundary Consistency” is:

Extra Condition. Let τ index future periods, then $\forall (m, J_i) \in B(\mathcal{W})$

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

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

Testing for Boundary Consistency.

Fix $(m, J_i) \in B(i)$. Simulate independent sample paths from it with initial J_{-i} drawn from the empirical distribution of $p^e(J_{-i}|J_i)$. Calculate mean, $\hat{W}(m|J_i)$, and the variance, $\hat{V}(\hat{W}(m|J_i))$, of simulated sample path for each $(m, J_i) \in B(i)$.

Basis of Test. Average

$$\frac{\left((W(m^*|J_i) - \hat{W}(m|J_i))_- \right)^2}{\hat{V}(\hat{W}(m|J_i))}$$

over $(m, J_i) \in B(i)$ and then a weighted average of these over boundary points. This is an Inequalities based test and one needs to simulate the test statistic's critical values.

Each path which we simulate either will or will not re-enter \mathcal{R} . Provided prior test is satisfied we have the correct expectation of the future value from any $(J_i, J_{-i}) = s \in \mathcal{R}$.

Let; r index simulation samples,
 γ_r index the periods simulated for sample r ,
 γ_r^* be the first period when $s_{\gamma_r} \in \mathcal{R}$ (or some sufficiently large number if it does not enter),
 $\{s_{\gamma_r}\}_{\gamma=1}^{\gamma_r^*}$ be the sequence of states simulated for sample path r .

Then an unbiased estimate of the actual value of the feasible play is

$$\hat{W}_r(m|J_i) \equiv$$

$$\sum_{\gamma_r=1}^{\gamma_r^*-1} \delta^{\gamma_r} \pi(m(J_{i,\gamma_r}), m^*(J_{-i,\gamma_r})) + \delta^{\gamma_r^*} W(m^*|J_{i,\gamma_r^*}).$$

If there are R simulated paths, let $\bar{W}^R(m|J_i)$ be their average, and $Var[\bar{W}^R(m|J_i)]$ be the standard estimate of the variance of this average.

Let $B(J_i) = \{m : (m, J_i) \in B\}$ and $\#B(J_i)$ be the number of elements in $B(J_i)$. So

$$T(J_i) = \frac{1}{\#B(J_i)} \sum_{m \in B(J_i)} \left(\frac{[\overline{W}^R(m|J_i) - W(m^*|J_i)]_+}{W(m^*(J_i))} \right),$$

is a measure of the deviation of the boundary point from boundary consistency.

Let $\mathcal{J}_B = \{J_i : (b, J_i) \in B \text{ for at least one } b\}$, $h(J_i)$ be the number of times times the point J_i was visited in the test run, and

$$q(J_i) = \frac{\{J_i \in B\}h(J_i)}{\sum_{J_i \in B} h(J_i)}.$$

Then our test statistic is

$$T(B) = \sum_{J_i \in \mathcal{J}_B} q(J_i)T(J_i).$$

We have to simulate its distribution under the null that $W(m|J_i) = W(m^*|J_i)$ for each $(m, J_i) \in B$ (this insures the size of the test), and check

whether the 95th percentile of the simulated distribution is larger than $T(B)$. We accept

H_0 : Boundary Consistency

if and only if it is not.

Maintenance in an Electricity Market

Full revelation every T periods guarantees existence of "standard" notion of equilibrium.

Computational Questions.

- Does the impact of increasing the number of periods of memory before full revelation taper off?

- Do reasonable bounds on memory lead to policies which approximate unbounded policies?

Compare:

(i) AsI with full revelation (T periods),

(ii) to AsI with T periods of memory but no revelation,

(iii) to AsI remembering last time each generator shutdown ($\leq T$ periods).

Note. Last two do not satisfy equilibrium conditions. However they do ask whether aggregating states in this fashion can mimic computational results.

Electric Utility Questions.

- Compare AsI to Social Planner (computed via standard contraction mapping). This is a social planner which is constrained to bid the same functions as the AsI does, and earns its return to capital through the difference between the price and its marginal cost (wanted to emphasize difference in maintenance behavior). The planner maximizes the sum of consumer and producer surplus.
- Compare AsI to FI.
- Compare results for different primitives: in particular we found rather striking differences in outcomes when the ratio of generation capacity to demand changed.

Model Details.

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

$$\begin{aligned} \pi_{i,t} &= \sum_r p_{t,r} q_{i,r,t} - \sum_r c_i(\omega_{i,r,t}, q_{i,r,t}) - f_i \sum_r \{m_{i,r,t} = 1\} \\ &\equiv \pi_i(\omega_i, m_i, b(m_{-i}), d) \end{aligned}$$

$$\begin{aligned} m_{i,r,t} = 0 &\Rightarrow \omega_{i,r,t+1} = \omega_{i,r,t}, \\ m_{i,r,t} = 1 &\Rightarrow \omega_{i,r,t+1} = \bar{\omega}_{i,r} \text{ (}\bar{\omega}\text{=restart state)}, \\ m_{i,r,t} = 2 &\Rightarrow \omega_{i,r,t+1} = \omega_{i,r,t} - \eta_{i,r,t} \\ &\text{with } P(\eta) > 0 \text{ for } \eta \in \{0, 1\}. \end{aligned}$$

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).

Parameters

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

Initial conditions. Set high \Rightarrow experimentation.

$\pi_i^{E,k=0}(m_i, J_i) = \pi_i(m_i, m_{-i} = 0, d, \omega_i)$, and

$W^{k=0}(\eta_i, m_i | J_i) = \frac{\pi_i(m_i, m_{-i}=0, d, \omega_i + \eta_i(m_i))}{1-\beta}$.

Note: Whenever we have used initial conditions this high we have satisfied boundary consistency conditions. But there is a cost in terms of compute time.

Convergence. 500 million iterations $\Rightarrow \mathcal{L}^2(\mathcal{P}(\mathcal{R})) \geq .99995$. $\mathcal{L}^2(\mathcal{P}(\mathcal{R})) \geq .99$ at 200 and flattens out at minimum between 250 and 350 million.

Computational Results.

- **Approximation from reducing the state space.** As long as you have a rich enough partition of the state space the approximations do well. Approximate $T = 6$ well with $T = 5$ (but

not $T = 4$), and approximate full revelation $T = 5$ well with finite history s (but not finite history τ).

- **Compute times.** Largely determined by burden of finding points in memory (storage: public information by tree structure, and private information with a hash table conditional on public info). So it varies (almost linearly) with the size of recurrent class.

- **Size of Recurrent classes** *Quite large* for simple model ($\approx 250k$). For richer models this becomes problematic;

- (i) hard for us to analyze, and

- (ii) harder to believe that agents actually condition on this much information.

It becomes noticeably small when we use aggregators, with consequent fall in compute time.

- Might look for ways (including testing) to reduce memory requirements. Likely to be problem specific.

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: Shutdown 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.