

Multiple Equilibria and Selection by Learning in an Applied Setting*

Robin S. Lee[†] Ariel Pakes[‡]

February 2, 2009

Abstract

We explore two complementary approaches to counterfactual analysis in an applied setting where multiple equilibria are possible. The setting is an ATM network example empirically investigated in Ishii (2005), and our counterfactual examines the reallocation of ATMs among banks following a hypothetical merger accompanied by an unexpected cost shock. The analysis conditions on Ishii’s estimates of the remaining parameters and on the observed “initial” conditions in the market. First we simply enumerate the possible equilibria. Then we examine how different learning algorithms “select out” among them. We find that there are only a small number of equilibria that are consistent with the observed market conditions and that their similarity and proximity to one another allows the researcher to bound counterfactual outcomes with some degree of confidence. The learning process induces a *distribution* of limiting play across these equilibria, and that distribution is sensitive to the specification of both the cost shock and of the learning process. If we are to use learning algorithms to improve on the bounds obtained from enumeration, evidence on the relevance of alternative learning processes would be helpful.

1 Introduction

One reason for obtaining detailed empirical estimates of the primitives underlying market outcomes is to enable a realistic analysis of the likely impact of policy and environmental changes. The

*We are grateful to Joy Ishii for supplying the necessary data.

[†]NYU Stern and Yahoo! Research.

[‡]NBER and Harvard University.

analysis of these changes, or “counterfactuals,” when there are multiple possible equilibria poses a problem to applied researchers: the model will not generate a unique counterfactual prediction. If the change in primitives moves the market to a state not observed in the past, the available data can not provide an answer to the choice among possible equilibria. Indeed, even if the counterfactual moves the market to a state that has been observed and assumptions which allow us to use the data to identify which of the possible equilibria were played in the past are viewed as acceptable, we can not rule out the possibility that the change in primitives will change the equilibrium selection mechanism. Moreover the appropriate selection mechanism is likely to depend on historical events the researcher does not have information on.¹

This paper explores two potentially complementary approaches to addressing the problem of multiple equilibria in counterfactual analysis. We do so in the context of a particular application; a network example drawn from Ishii (2005). The example is a simultaneous move game whereby competing banks within a local market choose the number of ATM machines to offer, with payoffs to the game generated from Ishii’s estimated structural model. We propose a merger of two banks which is accompanied by an unexpected cost shock of running an ATM, and analyze the reallocation of ATMs among the remaining banks. There are multiple potential equilibria of this game.

We begin by simply enumerating all the potential equilibria. Although there are nearly 200,000 possible allocations, we find that at most three of them constitute an equilibria for the specifications of primitives we investigate. The equilibria generated by a given specification are very similar to one another in that the number of ATMs operated by each bank differ across equilibria by at most one. Further the “comparative static” relationships between the sets of equilibria when we vary the specification make economic sense – e.g., if an allocation becomes a new equilibrium when costs of running an ATM are lower and another allocation no longer is an equilibrium, the new equilibrium always has a larger total number of ATMs than the other. This indicates that, when feasible, enumeration may lead to useful bounds on possible counterfactual results.

Next we explore the implications of using different learning algorithms to “select out” equilibria. That is, we model the process by which agents adjust over time to changes in their environment, and then follow those adjustments until no agents wishes to deviate from their chosen strategies. Though much theory has been written on learning models and their properties (c.f. Fudenberg and Levine

¹Accordingly few papers have explicitly tackled the issue of multiplicity in applied counterfactual analysis. A notable exception is Jia (2008), who analyzes a supermodular game between two agents and, by exploiting the result that equilibria in this class of games lie in a lattice (c.f. Topkis (1998)), is able to compute counterfactual results at the extremum equilibria to bound results. Allcott (2009) utilizes learning processes similar to those suggested here to select among multiple equilibria in his analysis of wholesale electricity pricing.

(1998), Young (2004), and the literature cited therein), there has been little if any experience with applying them to choose among equilibria. We examine both a best-response dynamic whereby agents choose the best reaction to the actions of its competitors in the previous period, and a fictitious play variant whereby agents instead play a best response to the distribution of previous play by competitors. We find that the variance in the cost shocks can cause a *distribution* of rest points from a given initial condition, and that distribution has a notable dependence on both the cost specification *and* on the learning process. Thus if we are to use learning algorithms to improve on the bounds obtained from enumeration, it would be helpful to obtain some evidence on which learning processes are more relevant in which environments.

2 A Network Example

We focus on the competition between banks for consumers via ATM networks studied in Ishii (2005). Ishii analyzes the choice of ATMs by banks and the impact of those choices on consumer banking decisions and equilibrium interest rates. In the process she provides detailed estimates of the primitives from a two period model which allows her to construct the profits of each bank conditional on *any* set of ATM networks for all banks in a number of markets. We take the information she provides on Pittsfield, Massachusetts and analyze the likely impact of a change in Pittsfield’s banking environment.

For the change we posit a hypothetical merger accompanied by an unexpected shock to Pittsfield’s economy which changes the costs of running an ATM. There were eight banks before the merger, so we examine the actions of the seven remaining banks in the market. We assume the merged bank has a profit function which consists of the sum of the profits from the two banks which merged, and begins with their combined number of ATMs. This provides us with an “initial” allocation of ATMs across the seven banks given by the vector $(9, 0, 3, 1, 0, 0, 1)$.² We note that, as is often the case in empirical work, there is significant heterogeneity across the firms inherited from past actions and events. In particular the banks differ in the number and locations of their branches and in the amenities they provide customers. We are assuming that these characteristics of the banks *do not* change.

²We focus only on “out-of-branch” or stand-alone ATMs; thus, a bank with 0 ATMs may still have these machines within one of its staffed branches.

The realized costs of agent i if it uses n_i ATMs in period t are given by:

$$C(n_i, t) = [b_{0,i} + b_{1,i,t}]n_i + b_2n_i^2 \quad (1)$$

where $(b_{0,i}, b_2)$ are known constants and $b_{1,i,t}$ is cost shock specific to firm i and time t .³ We assume these cost shocks to be i.i.d. draws from a normal distribution with mean μ and variance σ^2 that is common across firms. We assume μ is initially unknown. For simplicity, we also assume switching costs and fixed costs of each machine to be 0; we only focus on the per-period operational costs given by (1).

Firms do not know their future cost shocks before they chose the number of ATMs they operate in the next period, and we focus on Nash Equilibria in expected costs. In the first period after the merger, each firm receives its own realization of the cost shock $b_{1,i,t}$. As firms realize their costs, they use an average over previous post-merger cost draws to form their expectation of costs for the next period (μ). There are no dynamics other than that induced by learning about the likely value of the cost shocks and the likely play of competing firms.

Before proceeding, we emphasize that all our following results *condition* on the significant heterogeneity across the firms inherited from past actions and events, and the heterogeneity across profit functions typically found in actual data sets (including our own).

3 Number and Nature of Equilibria

The first part of the analysis proceeds by simply enumerating the “limiting equilibria”: i.e., the Nash equilibria when all firms know the expected value of the cost shock. Since banks are asymmetric, there are 170,544 different allocations of up to 15 ATMs among seven banks.⁴ Table 1 lists all equilibrium allocations when firms know the expected value of the per-period cost shock. We report results for different values of μ , which we set to be (\$20K, \$15K, \$10K, \$0).

The initial post merger allocation is (9,0,3,1,0,0,1), and does not constitute a Nash equilibrium for any of our cost specifications. Although the number of equilibria depends on the exact specification, it is always *strikingly small* in comparison to the number of total possible allocations. In two specifications the number of equilibria is three, in one specification there is two, and when

³In our example, we set $b_2 = \$20K$ and choose firm-specific $b_{0,i}$ values so that pre-merger allocations are an equilibrium. We assume the merged firm obtains the lowest value of $b_{0,i}$ among the two merging firms. These parameters generate average per-period costs of \$85K per ATM at the initial post-merger allocation.

⁴The number of ways of allocating at most n objects among m different agents is $\sum_{k=0}^n \binom{m+k-1}{k}$.

Table 1: **Possible Equilibria for Four Mean Cost Specifications**

Mean Cost (μ)		20,000	15,000	10,000	0
ATM Allocation	# of ATMs	Is Allocation An Equilibrium?			
(4,0,4,0,0,1,1)	10	Yes	No	No	No
(5,0,3,0,0,1,1)	10	Yes	No	No	No
(4,0,4,0,0,1,2)	11	No	Yes	No	No
(4,0,4,0,1,1,1)	11	No	Yes	Yes	No
(5,0,3,0,1,1,2)	12	Yes	Yes	Yes	Yes

the cost shock has mean zero, there is a unique equilibrium. Indeed, there are only five different allocations that are equilibria in at least one of our four cost specifications.

Within a specification for costs, the different equilibria are quite similar to each other. There are no two equilibria for the same cost specification in which one firm differs in its number of ATMs by more than one ATM, and the maximum difference in total number of ATMs across equilibria for a given cost specification is two; no matter costs and no matter the equilibria chosen, the first firm is always the largest firm and always has either 4 or 5 ATMs, and similarly the second and fourth firms never find it profitable to have any ATMs.

Across cost specifications, no firm changes its equilibrium number of ATMs by more than one. Also “comparative static” results on the relationship between equilibrium sets seem to make economic sense: if an allocation which had been an equilibrium is no longer an equilibrium when we lower the cost, the high cost equilibrium that drops out is always the equilibrium with the least number of ATMs; and if an allocation becomes an equilibrium allocation when it had not been one at the higher cost, the new equilibrium allocation always has a larger total number of ATMs than the equilibria that are dropped out, and those that are dropped are always the equilibria with the lowest number of ATMs.

4 Equilibrium Selection by Learning

The second part of our analysis examines the implications of assuming that different learning processes govern the adjustment of firms to changes in their environment. Here, the changes comprise a merger accompanied by a cost shock. Motivated by our application to modeling a

counterfactual, we assume firms begin in the initial post merger allocation of ATMs, and adjust from there.

We assume that each firm expects its costs to be drawn from a distribution whose mean was equal to the sample mean of the cost shocks it had received since the regime change. For each specification of costs, we consider two different learning processes which differ in each firm’s beliefs about its competitors’ play. In the first process, each firm believes its competitors’ will play the same strategy in the current period as they did in the prior period, and thus each firm will play its best response to last period’s play (*best response dynamic*). The second process assumes that each firm believes the next play of its competitors will be a random draw from the set of tuples of plays observed since the regime change, and chooses its optimal strategy appropriately (*fictitious play*). To see how these learning processes react to noise, we also tried three different specifications for the coefficient of variation (CV) of the cost shock: (25%, 50%, 100%). This works out to be 6 variants for each value of $\mu \in \{\$20K, \$15K, \$10K, \$0\}$, or 24 total cases.

The possible equilibria for each base specification when expected costs equal the true mean costs are those described before in Table 1. We compute 1,000 trial runs for the best response dynamic and 100 runs for fictitious play in each specification. Each run is stopped when we have converged to a single allocation, where convergence is defined as having remained in the same state for 50 iterations. This location was viewed as a “rest point” of the process. Note that *all* rest points are Nash equilibria of the game where each agent knows its mean costs.⁵ Table 2 provides the fraction of rest points at various equilibria for the different cost specifications. With the exception of less than 2% of runs under one specification with the best-response dynamic, there were no problems with convergence.⁶

The variance in the cost shocks can cause a *distribution* of rest points, and this distribution depends on the underlying specification of the shock. This dependence of the distribution of equilibria on the learning process is disconcerting because of the lack of empirical evidence on the relevance of alternative learning models. On the brighter side, the distribution of the number of ATMs from the lower cost specifications always stochastically dominated those from the higher cost specifications. Also in these results and others not reported here, when we increased the coefficient of variation of the cost shocks the dispersion of rest points seemed to increase in both the best reply

⁵Since agents’ expected costs eventually converge to the true mean, any rest-point must to the best response dynamic must satisfy the Nash conditions. This is also true under the fictitious play dynamic (c.f. Fudenberg and Levine (1998)).

⁶When there was a problem with convergence, the routine cycled between two allocations where two firms would alternate between adding and dropping or dropping and adding an ATM in each period.

Table 2: **Fraction of Rest Points at Alternative Equilibria**

Mean (μ)	20,000			15,000			10,000			0		
CV (σ/μ) ^a	1	.5	.25	1	.5	.25	1	.5	.25	1	.5	.25
Best Reply												
4040011	.89	.87	.82									
5030011	.10	.10	.13									
4040012				.27	.14	.01						
4040111				.40	.21	.02	.04 ^b	.00	.00			
5030112	.01	.03	.06	.33	.65	.97	.94	1.0	1.0	1.0	1.0	1.0
Fictitious Play.												
4040011	.47	.41	.41									
5030011	.34	.44	.30									
4040012				.00	.00	.00						
4040111				.10	.01	.00	.00	.00	.00			
5030112	.15	.15	.29	.90	.99	1.0	1.0	1.0	1.0	1.0	1.0	1.0

The initial condition is (9,0,3,1,0,0,1) for all runs and is never an equilibrium based on true expected costs.

^a CV is the coefficient of variation of the cost shock. For the base specification where $\mu = 0$, the variance of the cost shocks were set to be the same as when $\mu = 20,000$.

^b In this specification under Best Reply, approximately 2% of trials resulted in “cycling.”

and fictitious play dynamic, and it tended to take more iterations before we satisfied our stopping rule. Finally, there appears to be a sense in which certain equilibria are more “dominant” than others: for the lower CV specifications, certain equilibria are played in the majority of runs (e.g., the 12 ATM equilibrium for $\mu = 15K$ and $10K$), and certain equilibria also seem to be reached very infrequently (e.g., the 11 ATM equilibria).⁷

⁷Our experiments generated one other result of note. There was a tendency for the fictitious play dynamic to generate rest points which, in terms of the number of ATMs, stochastically dominated the equilibria when a best reply dynamic was used. This was not totally a result of the twin facts that fictitious play puts more weight on history and we started at an initial condition in which more ATMs were serviced; the same result was evident, though to a lesser degree, when we started at alternative initial conditions with fewer ATMs.

5 Concluding Remarks

We considered two approaches for analyzing counterfactuals in a particular applied problem in which multiple equilibria are possible. The first simply enumerated the total number of equilibria and examined their relationship to one another. To the extent that our findings here are indicative of what might happen in other applied problems, they are good news for the applied researcher. The small number of equilibria should make it feasible to do the enumeration, and the fact that the equilibria for a given cost specification are similar to each other is likely to imply that they have implications which are also similar. Finally the expected comparative static results (on sets) hold when we compare across the equilibria generated by the different cost specifications.

As noted these results are likely to rely on the fact that the actual profit functions have a substantial amount of heterogeneity built into them. For example were we to eliminate this “history” and assume that banks chose the number and location of their branches along with the number of their ATMs, or even just allow them to trade the branches at the current locations, the results would no doubt be quite different.⁸ On the other hand the applied problems we typically analyze come with a certain amount of inherited heterogeneity and restrictions to change built into them. So problems like the one analyzed here may not be uncommon. Indeed the applied researcher would find a characterization of profit (or value) functions that generate a small number of (possibly “well behaved”) Nash equilibria extremely useful.⁹

The second approach we examined consisted of utilizing different learning algorithms to “select out” equilibria. The results on these experiments showed that the distribution over equilibria we obtained depended not only on the particular cost specification, but also on the particular learning process. We also found that certain equilibria appeared to be more “dominant” than others and some equilibria were very infrequently played. These results indicate that evidence on the relevance of alternative learning processes, and/or a characterization of the “attractiveness” of different equilibria, would help to improve on the bounds generated by enumeration.

References

ALLCOTT, H. (2009): “Real Time Pricing and Electricity Markets,” mimeo.

⁸It is worth noting that asymmetry alone helps reduce the number of potential equilibria: if firms were symmetric and $(4, 0, 4, 0, 0, 1, 1)$ was an equilibria, then there would have to be at least $C(7, 4) \times C(4, 2) = 210$ different equilibria.

⁹Related work in the theory literature comprises, e.g., McLennan (2005) who analysis of the number of possible Nash equilibria for general payoff functions, and the analysis of the impact of different forms of heterogeneity in global games (c.f. Morris and Shin (2003)).

- FUDENBERG, D., AND D. K. LEVINE (1998): *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- ISHII, J. (2005): “Interconnection Pricing, Compatibility, and Investment in Network Industries: ATM Networks in the Banking Industry,” mimeo.
- JIA, P. (2008): “What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Industry,” *Econometrica*, forthcoming.
- MCLENNAN, A. (2005): “The Expected Number of Nash Equilibria of a Normal Form Game,” *Econometrica*, 73, 141–174.
- MORRIS, S., AND H. S. SHIN (2003): “Global Games: Theory and Applications,” in *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, ed. by M. Dewatripont, L. P. Hansen, and S. J. Turnovsky, vol. 1, pp. 56–114. Cambridge University Press, Cambridge, U.K.
- TOPKIS, D. (1998): *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ.
- YOUNG, H. P. (2004): *Strategic Learning and its Limits*. Oxford University Press, Oxford, UK.