Implications of the Small Gain Theorem in the Design of an Economic Laboratory

N. Tran, C. Giraud-Carrier, K. Seppi, and S. Warnick

Information Dynamics and Intelligent Systems Group
Department of Computer Science, Brigham Young University, Provo, UT 84602

***Abstract*— This report explores the challenges and opportunities inherent in developing a retail laboratory. In particular, retail is considered as a feedback phenomenon, and some fundamental limits on learning from data are identified. The small gain theorem is then reviewed, and the unique value of a laboratory enabling controlled, live experimentation is highlighted in light of the theorem's implications.**

I. Introduction

One of the difficulties in studying economics is the absence of a true laboratory enabling controlled experimentation. Unlike physics or chemistry, economics has been compared to astronomy as an observational, rather than an experimental, science.

Nevertheless, there are significant differences between economics and observational sciences like astronomy. While there is little astronomers can do to affect the astrodynamics they observe, entrepreneurs, mechandisers, and other strategic and tactical managers make critical economic decisions every day that impact the life and death of firms and drastically affect the economic landscape. Certainly market dynamics are not as out of reach as one might expect from an observational science.

Moreover, economics retains a sense of quantification more like physics than other behavioral sciences such as psychology. The fact that units of exchange are well defined enable the objective interpretation of economic behavior, and firms are often well instrumented to count and record more data than they ever actually use.

Another distinction between economics and non-experimental sciences include the fact that the economic systems we engage are engineered, not natural, phenomena. This observation lead Hal Varian to view economics as a policy science [5], more like engineering or medicine than physics or biology, and may explain the abundance of economic workshops (enterprizes) and dearth of laboratories.

Noting these differences between economics and non-experimental sciences, certainly an economic laboratory seems feasible, at least on the surface. In fact, with data analytics increasingly supporting business decisions, firms seem to be attempting to create in-house laboratories measuring the impact of their decisions and trying to better understand the nature of their market environment.

This work is supported by grants from the Rollins Center for eBusiness, the Office of Research and Creative Activities at BYU, and the BYU Bookstore. Please direct all comments and questions to Sean Warnick at `sean@cs.byu.edu`.

The next section offers some background and motivation in the development of laboratory-like environments in economics, and Section III discusses specific challenges in the creation of such a laboratory. Section IV then introduces retail as a feedback phenomenon and identifies some fundamental limits on learning from data. Section V then reviews the small gain theorem and presents an argument for the unique value of a laboratory enabling controlled, live experimentation in light of the theorem's implications.

II. Background

In his article, "Stochastic Control for Economic Models: Past, Present, and Paths Ahead," [3] Kendrick reports how in the 1950s a group of four bright, young economists at Carnegie-Mellon (Charles Holt, Franco Modigliani, John Muth, and Herbert Simon) leveraged control-theoretic techniques in microeconomics by computing variables for production, inventories, and the labor force in a firm. Notably, however, these scientists were not merely satisfied to develop the mathematics supporting their ideas, but they sought to measure their ideas against a testbed, which they found in a local paint factory that was willing to supply them with data.

More recently, the field of Revenue Management (RM) has undertaken efforts to design intelligent offers for firms ever since the Airline Deregulation Act of 1978 [4]. Although initially focused on the Airline Industry, RM has expanded its focus to include retail, driven strongly by the explosion of scanner data and the dawn of internet shopping. A central theme in all of these efforts is the connection of theory to practice. Often, parameters of a mathematical model are fit from industry or a firm's data, and the subsequent optimal or otherwise defendably-good policy is then derived and reported. Likewise, other communities, such as the scientific marketing, machine learning, data mining, and personalization communities, etc. have sought methods to help firms intelligently design their market interactions. Sometimes the ideas are shown to be consistent with real data, and, sometimes, an idea is implemented.

Nevertheless, in all of these cases, details of the implementation are often proprietary and unavailable for scientific critique. As a result, a gap exists between reported results and the beliefs of the scientific community (and many savvy managers). Scientists argue that much of the work is hyped, that it often implies the generalizability of results that may be very specific to a particular situation, while Practitioners claim that the more scientific work is overly simplified and thus irrelevant to their real-world problems. In spite of all the progress and work towards an experimental economics that can translate data into specific help on key decisions, the need for an open laboratory, where scientific methods can be tested and compared, remains crucial.

III. A Retail Laboratory

The essential feature of a laboratory, in this sense, is the ability to conduct controlled experiments and measure the results. Firms that implement academic work (whether in-house research or otherwise) and measure the results take a huge step in this direction. Nevertheless, since no firm can ”reset” the internal variables of its market at will, there is a question of repeatability that makes it difficult to interpret experimental results. For example, is the increase in the number of transactions a firm observes today the result of the experimental price change it made yesterday, or were customers already in the process of warming up to the firm–and today they finally made a purchase without even knowing about yesterday's change in price? With the ability to completely control experiments, scientists could ”reset” people's attitudes toward the firm (and all other market factors, such as competition, etc.) and simply compare behavior's resulting from different stimuli. Clearly, however, this is not possible in the retail setting, and it has been a significant challenge to the idea of any kind of economic laboratory.

To overcome this challenge, scientists need a method for clearing up the ambiguity introduced by having no direct control over internal market variables or other market forces. That is, scientists need a way to observe behavior caused by the stimuli they apply to the system, independent of internal market variables or market forces over which they have no control nor the ability to observe.

IV. Retail as a Feedback Process

We begin our discussion with the idea of a firm. Our firm will begin its life with a fixed amount of capital $C(0)$ and a list of n wholesale goods it can choose to purchase. There may be many potential suppliers of these goods, so the actual information the firm receives is a matrix of wholesale prices (effective costs) we denote as a $p \times n$ matrix $P_W(k)$, where $k = 0, 1, 2, ...$ is the number of days since the firm began operation. Thus, $P_W(k)$ are the prices the firm receives on the morning of day k, and we assume they will not change until the next morning.

The firm purchases wholesale goods on day k at price $P_W(k)$ by specifying a matrix $X_W(k)$, which identifies how many of the n goods from each of the p suppliers the firm will purchase at the current price, $P_W(k)$. Of course, the firm has limits on its spending characterized by the constraint that $\sum_{i,j} x_{w(i,j)} p_{w(i,j)} \leq C(k)$, where $x_{w(i,j)}$ and $p_{w(i,j)}$ are the (i,j)th elements of $X_W(k)$ and $P_W(k)$, respectively.

The firm then makes offers to its retail market for a price of its choosing. Since the firm may offer different effective prices to its m customers (or different store locations, or other type of channel), say through coupons or special offers, etc., the pricing decision of the firm is represented by a $m \times n$ matrix $P_R(k)$, which means the prices available to the m retail customers (or channels, or segments) of the n goods on the morning of day k. Note that goods that the firm has not purchased wholesale are listed as available for very high prices.

The retail market then responds with purchases throughout day k characterized by a matrix $X_R(k)$. This $m \times n$ matrix describes how much of each of the n goods each of the m retail customers or segments purchased that day. At this point, money from the day's transactions are collected, and the next morning's available capital becomes $C(k+1) = C(k) - \sum_{i,j} x_{w(i,j)} p_{w(i,j)} + \sum_{i,j} x_{R(i,j)} p_{R(i,j)}$.

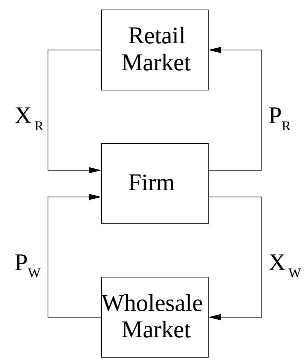

Fig. 1. A retail firm interacts with wholesale and retail markets.

Although real prices are quantized by the fundamental monetary unit, and the unavailability of a good may be considered ontologically different than associating a very high price with the good, we will ignore these details for simplicity and consider that $P_W(k) \in \mathbb{R}^{p \times n}$ and $P_R(k) \in \mathbb{R}^{m \times n}$. We thus define the firm as an operator $F : (\mathbb{R}^{m \times n} \times \mathbb{R}^{p \times n}) \rightarrow (\mathbb{R}^{m \times n} \times \mathbb{R}^{p \times n})$, its wholesale market as an operator $M_W : \mathbb{R}^{p \times n} \rightarrow \mathbb{R}^{p \times n}$, and its retail market as an operator $M_R : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$, which yields the picture shown in Figure 1. More generally, we could consider cases where the wholesale and retail markets are interrelated, yielding a combined market operator $M : (\mathbb{R}^{m \times n} \times \mathbb{R}^{p \times n}) \rightarrow (\mathbb{R}^{m \times n} \times \mathbb{R}^{p \times n})$. Finally, instead of thinking of these operators as mapping decisions at each time instant k, we can think of them as mappings between function spaces defined over $\mathbb{R}^{m \times n}$ and $\mathbb{R}^{m \times n}$. Modeling a situation where every price and purchase is finite on any particular day suggests we characterize the vector space of the firm's actions as $\mathcal{U} = l_\infty^{(m \times n) \times (p \times n)}$ and observations as $\mathcal{Y} = l_\infty^{(m \times n) \times (p \times n)}$. Here, the notation l_∞ refers to the space of infinite sequences of bounded amplitude, and the superscript $l_\infty^{(m \times n) \times (p \times n)}$ merely indicates a (pair) of matrices of such sequences. This results in a very general feedback structure as shown in Figure 2, where C is an internal state of the firm. We assume that such interconnection is *well-posed*, in that solutions of the system $u(k)$ and $y(k)$ exist and are unique.

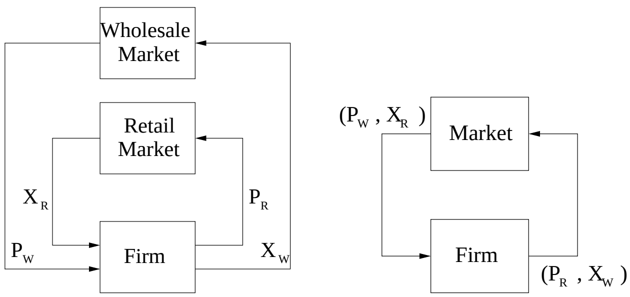

Fig. 2. By rearranging Figure 1 (left), we note that the distinction between a retail and wholes market simply imposes structure on the more general feedback relationship (right).

Notice that in this view of retail, a firm's interaction with the market is to buy and sell. Coupons, markdowns, special promotions, advertising, etc. are all modeled as equivalent modes of communicating offers to the firm's customers. Likewise, all modes of communication with a firm's suppliers are considered equivalent. Although this may not be precise for real economic systems, it is a reasonable simplification for the purposes of our analysis.

Furthermore, note that usual classifications of the firm as a price-setter or price-taker, or of the market as being competitive, oligopolistic, or monopolistic are unnecessary here. The firm makes its procurement and pricing decisions with the best information it has at any particular time, and information about whether it is operating in a competitive environment or has some special market power will be reflected in how the market responds.

With this view of the feedback process, where all market influences and disturbances have been lumped together into one operator $M : \mathcal{U} \to \mathcal{Y}$, we want to explore fundamental limits on validating a learning process. We may suppose that M has a state space representation of the form

$$\begin{aligned} x(k+1) &= f(x(k), u(k), k) \quad x(0) = x_o \\ y(k) &= g(x(k), u(k), k) \end{aligned} \tag{1}$$

where u is a vector of size $p \times n \times m \times n$ (the number of distinct decisions the firm can make at time k), x is a presumably very large vector of internal market states (that we, nevertheless, assume to be finite), and y a vector of market observations (same length as u). We assume that the firm has no special information about its market, and thus only knows $[u(k), y(k)]$ over a finite duration $k = 0, 1, 2, ..., t$. In particular, the firm does not observe the internal market states x, nor does it have complete information of f or g.

Once in operation, however, an intelligent firm will begin to consider its observations $[u(k), y(k)]$, $k = 0, 1, ..., t$ and use them to develop some understanding of how its market behaves. It may accomplish this by hiring experts who retain some partial information of the market dynamics f and g, or by employing novel learning processes that analyze its data in effective ways. However the firm operates, though, it will develop a model or theory of how the market behaves and attempt to use this understanding to its advantage.

Let the firm's model or theory of the market be an operator $T : \mathcal{V} \subset \mathcal{U} \to \mathcal{Y}$. That is, the understanding a firm obtains about its market enables it to establish expectations about what it will observe for at least some of the actions it might take ($u \in \mathcal{V}$). Regardless of how the firm developed a particular theory T, however, it will need determine how well the theory approximates its true market before it will have the confidence to widely use it. Restricting the market to the situations where the theory applies ($\mathcal{V} \subset \mathcal{U}$), the firm would like to measure some notion of distance between M and T. For this purpose, we may define the gain of an operator to be

$$\|T\| = \sup_{u \neq 0} \frac{|Tu|_\infty}{|u|_\infty}. \tag{2}$$

This measure of operator size is especially suited to our choice of norms on the action and observation spaces ($\mathcal{U}$ and $\mathcal{Y}$) since it is induced by them through their norms $|\cdot|_\infty$ in the numerator and denominator of (2). That is, change the norms on $\mathcal{U}$ or $\mathcal{Y}$, and the induced norm $\|\cdot\|$ changes as well. Under this measure, then, the quality of approximation of a theory is found by computing $\|M - T\|$ over the set $u \in \mathcal{V}$ where the theory applies.

To evaluate the quality, in this sense, of any process for learning from data, then, one must compute $\|M - T\|$. This is difficult for a few reasons. First, since the operators are allowed, indeed, expected to be nonlinear, there is not an obvious simplification in the search over $\mathcal{V}$ for the action that produces the largest relative output. That is to say, to compute the supremum, effectively a search over all $u \neq 0 \in \mathcal{V}$ must be undertaken to find the input that yields the relatively largest output. Without some other structure to the problem, the situation could be as hopeless as mere trial and error. The trend in building terabyte-sized data warehouses in motivated, in part, by this sort of reasoning. Nevertheless, a complete search over all possible pairs $[u(k), y(k)]$ is impractical for any real firm to consider, and it is very difficult to ever know if you have searched "enough" without having searched everywhere.

Another difficulty in computing $\|M - T\|$ is the fact that in practice a firm can only observe the impact of a particular action over a finite duration. That is, the data pairs $[u(k), y(k)]$ are actually restricted to observations over finite intervals of time. This implies that to meaningfully measure the error between our theory T and the true market M, the error must be asymptotically stable, in that it will converge to an acceptable bound within a finite period and will never grow beyond it thereafter. In situations where we expect M itself to be asymptotically stable with a large enough rate of convergence, this may not pose a problem. Nevertheless, for any real situation where competition and other market forces continually impact the transactions observed by a firm, it seems unreasonable to assume that such strict stability will arise naturally. Thus, in situations where M is not asymptotically stable, *any learning algorithm yielding the model T and using a notion of approximation $\|M-T\|$ as a measure of its validity must have* perfect knowledge *of the unstable modes of M in order to cancel them and make $\|M-T\|$ computable.* This requirement for perfect information about the instabilities of the market is too much to ask, even from the most intelligent of learning processes.

V. The Small Gain Theorem

If the direct comparison $\|M-T\|$ is unreasonable, another approach may yet enable a computable measure of quality of a proposed market model. The method is based on the small gain theorem [1], [2]:

Theorem 5.1: Consider the well-posed feedback interconnection of finite-gain operators R and G, mapping between l_∞ spaces of compatible dimension. Then the closed-loop system is stable if $\|R\|\|G\| < 1$.

This theorem is helpful because the specific nature of the operators, such as their order (the number of hidden internal variables) or the complexity of their nonlinearities, etc. is irrelevant to the result. This allows us to use the theorem to deal with uncertainty and complexity in a feedback process.

In particular, the true market M can be factored into the feedback interconnection of a theory of the market T with the unknown remainder R. This feedback factorization often allows an unstable operator M to be decomposed into the feedback interconnection of two stable operators R and T. Thus, although the conditions of the small gain theorem require that the unknown part of the market, R, be stable, this is a weaker condition than demanding that the market itself, M, be stable.

The novelty of the approach arises when we consider actually implementing a specific policy F that translates observations from the market, y into specific actions u. Although F may be designed only with information of T, the small gain theorem says it will stabilize the real market M provided $\|R\|\|G\| < 1$, where G is the feedback interconnection of F and T. Note that we expect G to be stable since we design F such that it stabilizes T.

Thus, the scientist can test the quality of a given market theory T by designing a policy F that robustly and asymptotically stabilizes it with a desired rate of convergence. Analysis or simulation of this ”ideal” closed-loop system, G, then indicates how long it will take the feedback interconnection of F on any system M ”close” to T (in the sense that $\|R\|$ is small enough) to stabilize. This suggests a finite duration experiment that can be tested empirically by implementing the policy F ”live” on the true market M. Observing the closed-loop results over the required duration either results in success, meaning the closed-loop market is behaving as designed, or failure, which implies that the conditions of the theorem were not met and T is not close to M in that $\|R\| \geq 1/\|G\|$. These definite conclusions can be made, provided that a mechanism for implementing a closed-loop policy ”live” on the true market exists.

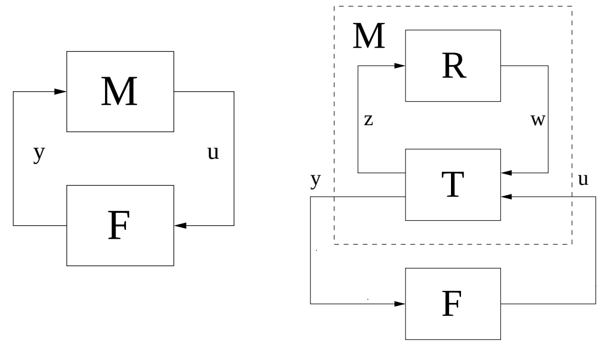

Fig. 3. The operator M can be factored as the feedback interconnection of R and T.

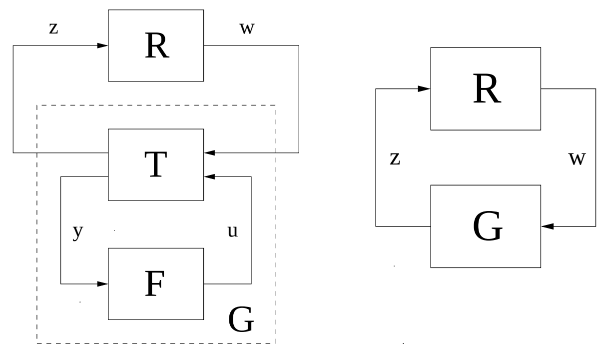

Fig. 4. Regrouping operators, the small gain theorem guarantees that F stabilizes M provided $\|G\|\|R\| < 1$.

This is the value of a true economic laboratory, where experiments can be conducted and observed, even without perfect control over the internal variables or outside influences of the market. Many validation processes use a database of previously recorded events, partition the database into a ”training” and ”validation” set, use the training set to generate a reasonable theory T, and then use the validation set to approximate something like $\|M-T\|$. Nevertheless, there are inherent limitations in this approach for real-world problems. An alternative ”closed-loop” scientific process leverages the small gain theorem and a laboratory enabling live experimentation to bypass these limitations and generate provably good policies for real-world problems.

References

[1] M.A. Dahleh and I.J. Diaz-Bobillo. *Control of Uncertain Systems: A Linear Programming Approach.* Prentice Hall, N.J., 1995.

[2] G. Dullerud and F. Paganini. *A Course in Robust Control Theory, A Convex Approach.* Springer, N.Y., 2000.

[3] D.A. Kendrick. Stochastic control for economic models: Past, present, and the paths ahead. *Ninth International Conference of the Society of Computational Economics*, Keynote Address, 2002.

[4] K. Talluir and G. Van Ryzin. *The Theory and Practice of Revenue Management.* Kluwer, Boston, 2004.

[5] H. Varian. What use is economic theory? *Prepared for the conference, ’Is Economics Becoming a Hard Science’*, Paris, France, 1992.