IEEE Access
The journal for rapid open access publishing

Received October 4, 2015, accepted November 9, 2015, date of publication December 17, 2015,
date of current version December 23, 2015.

Digital Object Identifier 10.1109/ACCESS.2015.2505730

Valuing Information in Complex Systems: An Integrated Analytical Approach to Achieve Optimal Performance in the Beer Distribution Game

KIMBERLY M. THOMPSON[1] AND NIMA D. BADIZADEGAN[2]
[1]College of Medicine, University of Central Florida, Orlando, FL 32827, USA
[2]California Institute of Technology, Pasadena, CA 91126, USA

Corresponding author: K. M. Thompson (kimt@aorm.com)

ABSTRACT Even seemingly simple systems can produce complex dynamics, which leads management professionals to develop tools for training, monitoring, and improving performance. Management simulators provide useful insights about human behavior and interactions, while computational and informational decision support tools offer opportunities to reduce inconsistencies, errors, and non-optimal human choices, particularly for complex systems that involve multiple decision makers, uncertainty, variability, and time. We use the context of a popular management simulator that teaches students about the bullwhip effect (i.e., the beer distribution game) to explore an integrated decision analytic, control theory, and system dynamics approach to the game that recognizes the value of available (imperfect) information and considers the value of perfect information to provide the optimal strategy. Using a discrete event simulation, we characterize optimal decisions and overall team scores for the situation of actual available information and perfect information. We describe our implementation of the strategy in the field to win the 2007 beer game world championship played at the 25th conference of the International System Dynamics Society. This paper seeks to demonstrate that better understanding of the system and use of available information leads to significantly lower expected costs than identified in prior studies. Understanding complex systems and using information optimally may increase system stability and significantly improve performance, in some cases even without better information than already available.

INDEX TERMS Value of information, beer distribution game, decision analysis, system dynamics, control theory.

I. INTRODUCTION

Even seemingly simple systems can produce complex dynamics [1]. The integration of multiple analytical tools may provide the best strategies for developing robust solutions in complex systems that involve multiple decision makers, particularly for decisions and games that involve uncertainty, variability, and time [2].

The large global economy and growing population depend on management of complex supply chains. Nonetheless, those responsible for managing supply chains remain largely invisible to most consumers, except in times of failure (i.e., shortages, gluts, major disruptions). The survival and growth of large companies depend on how they manage their supply chains, with companies like Wal-mart recognized for logistics leadership and their use of information [3] and innovations like cross-docking to minimize inventory holding times [4]. Major disruptions combined with poor management and poor use of information can lead to significant costs, and sometimes to enterprise failure [5].

The simplest supply chain involves one producer providing a single product exclusively sold by one retailer (i.e., a supply chain with 2 nodes), although most supply chains involve many more nodes and significant dynamic complexity. The complexity leads management professionals to develop tools to monitor and improve performance. Computational and informational decision support tools (including data collection and management systems and algorithms) offer opportunities to reduce the potential for inconsistencies, errors, and non-optimal human choices, and management simulators can provide useful insights about

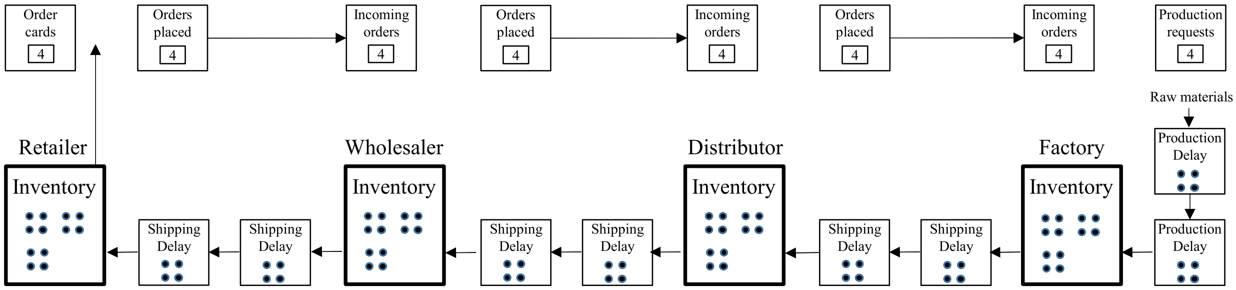

FIGURE 1. Schematic of the beer distribution game board and initial conditions for the four positions of retailer, wholesaler, distributor, and factory, the shipping and production delays, and the initial orders of 4 for the training rounds shown face up although the orders in the real game are face down.

human behavior and interactions to help train managers and students [6]–[8].

Over 50 years ago, Professor Jay Forrester created the beer distribution game and introduced the concept of the ‘‘bullwhip effect’’ [9]. The beer game is in the public domain, with materials and instructions available from the International System Dynamics Society [10]. Figure 1 shows the configuration of the game, which involves a supply chain with four players: the retailer, wholesaler, distributor, and factory [11]. Briefly, each player fills incoming orders from its inventory, which typically starts with 12 units, with the initial face-down ‘‘incoming orders’’ and ‘‘orders placed’’ slips all equal to 4 and a pre-determined deck of customer ‘‘order cards’’ for all rounds (i.e., simulated weeks) of game play placed for the retailer to draw one from the top for each round. The factory determines how much beer to produce each round, and shipping and production delays slow the movement of beer through the system. During the game, players can observe the full board, but not the face-down order cards or slips, such that only the retailer knows the customer order for the current round (and prior rounds) with certainty, and customer orders for future rounds remain uncertain until received. The game play involves training for the first 4 rounds, during which time the players receive orders for 4 units and instructions to always order 4 units, except at the end of round 4 when players can order whatever they choose. With the order slips face down (i.e., normal game play), at the end of the training rounds the board looks identical to the initial set up. Teams receive information about the costs of inventory and backlog (i.e., the inputs to the cost function) and instructions to seek the minimum total team costs (i.e., for the entire supply chain). At the end of each round (including the 4 training rounds), players record their current inventory or backlog (i.e., negative inventory representing the cumulative unfilled orders received). Each unit of inventory for each player held at the end of the round incurs a cost of $0.5, while each unit of backlog incurs a cost of $1, with the minimum cost per round of $0 assigned if the inventory and backlog equal zero. Customers waiting for orders continue to wait, such that backlog accumulates until filled. A seminal paper on the beer distribution game provides details of game play, documents the typical poor performance of actual teams with the game, and provides a systems dynamics ‘‘order-up-to’’ solution that offers a benchmark team score of $204 for the standard deck of order cards [11]. Subsequent studies [12]–[14] explored other strategies that led to similar benchmark team scores as low as $196 [13]. A recent study further highlights the key attributes of the beer game that make it a highly simplified supply chain, including ‘‘no random shocks in supply or demand, no capacity constraints on production, … no price variations, … no horizontal competition to trigger shortage gaming, and no cancellation of orders’’ [15].

In parallel over 50 years ago, decision analysts developed the concept of the value of information when faced with uncertain choices [16]–[19]. Numerous reviews demonstrate the benefits of consideration of the value of information in high-stakes decisions in a wide range of applications [20]–[24]. As the literature on supply chain management developed, an abundance of studies emerged related to characterizing the value of information in supply chains, including studies on information sharing and the bullwhip effect [25]–[28]. However, to date no prior studies considered an integrated decision analytic, system dynamics, and control theory approach to identify the optimal solution for the beer game considering the value of information.

In this analysis, we consider the incentives that motivate the decisions of the individual players and the potential optimal performance of the team from the perspective of decision analysis considering the value of information and the perspective of control theory assuming a single controller. We consider the information available during actual game play, such that all players can see the numbers of units in the different parts of the supply chain (if any), but they cannot communicate directly about orders, cancel orders, or observe the extent of backlog (i.e., negative units remain unobservable). Our analysis seeks to identify the optimal strategy for the beer game. The following section discusses the theory underlying our approach. We then discuss the methods we used to calculate the optimal orders and scores, our experience with field implementation of the optimal strategy, and our findings and insights.

II. THEORY

Recognizing the value of the imperfect information available and taking an integrated analytical approach leads to several observations. During each round, each player must decide whether to order 0 or some integer number of units from the upstream player on the right (or in the case of the factory how many units to produce), while remaining uncertain about the order that will arrive from the downstream player on the left with a one-round delay (or in the case of the retailer from the top card on the customer order deck). In the context of considering the value of perfect information, we begin by asking the question: what should each player order in the context of perfect information? This leads to the realization that orders placed by the retailer, wholesaler, and distributor represent the source of uncertainty for the wholesaler, distributor, and factory, respectively, and that the retailer, wholesaler, and distributor can place orders that maintain 0 inventory for the wholesaler, distributor, and factory, respectively. We then also recognize that given the unit costs and incentives, the location of inventory in the system does not matter.

Observation 1: From an inventory unit cost perspective, the location of any inventory in the supply chain does not matter.

Let the inventory held in the retailer for round r be $R[r]$, the wholesaler $W[r]$, the distributor $D[r]$ and the factory $F[r]$. The total inventory, $T[r]$, is given by: $T[r] = R[r] + W[r] + D[r] + F[r]$, with total inventory costs, $C[r] = 0.5T[r] + B[r]$, where $B[r]$ is the total system backlog. In the standard game play, inventory costs the same in every inventory position on the board while units of beer cost nothing while in shipping or production delays, so the desirable location of any inventory in the system does not depend on the unit cost of inventory. For a constant value of total inventory in the system, shifting inventory between any of the inventory positions does not affect the total cost of inventory for the team. ■

Theorem 1: Under a flat cost structure, all of the inventory in the system should be held in the retailer.

For any round r, let the retailer customer demand be $O[r]$. With $R[r]$ and $T[r]$ greater than zero at initialization of the game (i.e., when $r = 0$), the initial backlog in this system equals 0, and we can use the flat cost function from Observation 1. Two possible cases exist for values of $O[r]$ relative to $R[r]$:

1) $O[r] \leq R[r]$: $R[r+1]$ and $T[r+1]$ will both reduce by $O[r]$, without creating backlog. Cost for this round changes by $-0.5\,O[r]$.
2) $O[r] > R[r]$: $R[r+1]$ will go to 0, $T[r+1]$ reduces by $R[r]$ and creates $O[r] - R[r]$ in backlog. Cost for this round changes by $O[r] - 1.5R[r]$.

The change in cost for Case 1 does not depend on inventory distribution, but the change in cost for Case 2 does. We can minimize the cost for Case 2 by maximizing R, which occurs when $R = T$. This inventory distribution does not change for backlog $B[r] > 0$, because we can treat backlog as part of $O[r]$, and the same cases exist for $O[r] + B[r]$. ■

Given Observation 1 and Theorem 1, we observe that all inventory in the supply chain incurs costs from the time of production until exiting as part of a customer order filled by the retailer, except while moving in shipping and production delays. While the team players maintain responsibility for the orders that they place, the customer orders coming into the retailer remain exogenous and uncertain. Although the location of inventory in the supply chain does not matter from a unit cost of inventory perspective, insufficient inventory to fill a customer order in the retailer position creates unobservable and expensive backlog. Thus, moving inventory to the retailer makes it available for fulfilling uncertain customer demand.

Theorem 2: The cost-minimizing strategy for the retailer, wholesaler, and distributor is to place orders that achieve and maintain 0 inventory for the upstream players to their right (i.e., the wholesaler, distributor, and factory, respectively, see Figure 1).

Given Observation 1 and Theorem 1, the team should move all existing inventory in the system to the retailer, recognizing that inventory and backlog accumulate every turn. To achieve this, the retailer, wholesaler, and distributor should make ordering decisions each round that drop the inventory of the upstream player to 0 and maintain this 0 inventory, and they can do so independent of their incoming orders. By acting as a pass-through of goods from the factory to the retailer and effectively cross-docking all supply, these players simply move the inventory in the system to the point of use and they do not create unnecessary dynamics in the system. Thus, the optimal strategy recognizes that the retailer, wholesaler, and distributor can easily and rationally observe the inventory and amount of units in the shipping or production delay of the upstream player and determine their appropriate orders to maintain 0 inventory for the distributor, wholesaler, and factory. ■

For example, at the first opportunity (i.e., the end of the 4th training round), each of these three players should order 16 units for the standard game play (i.e., clear out the 12 units currently in the inventory of the upstream player plus the 4 units that will come in from the rightmost shipping delay or production delay of the upstream player prior to order filling). After placing an order that will clear any existing inventory of the upstream player, the retailer, wholesaler, and distributor should continue to order the amount in the rightmost shipping or production delay of the upstream player.

By incentivizing the retailer, wholesaler, and factory to minimize the cost of the inventory that their decisions actually control, these players each contribute the minimum possible to the overall team score (i.e., $0) for the upstream player after the training rounds. This occurs because the system of the retailer, wholesaler, and distributor represents a closed system from a control theory perspective, with one input and one externally-controlled output (Figure 2). These players cannot control the amount of incoming orders received, but completely control the inventory (and backlog) of their suppliers. Thus, these players control the location of the total inventory in the supply chain, and they can also generate

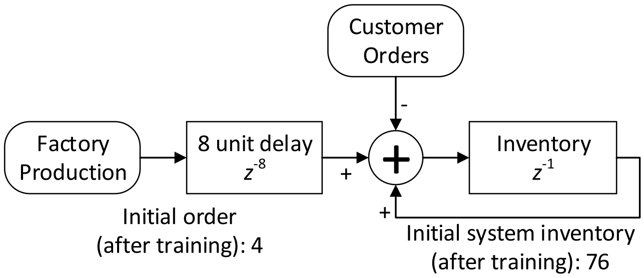

FIGURE 2. System block diagram showing system elements and their Z-transforms for the factory problem.

backlog for the factory, distributor, and wholesaler by over-ordering [15], but only the factory controls the total amount of inventory in the system. With these three players using a pass-through of *inventory* from factory to retailer strategy, from a control theory perspective, the factory, a single player, essentially controls an eighth-order linear system (Figure 2). We emphasize that this strategy contrasts with the concept of pass-through of *orders* upstream from the customer to the factory, which focuses on passing customer order information up the chain without any consideration of existing inventory downstream in the supply chain or the consequences of delays.

Observation 2: The training phase of the game implies fixed costs of $96.

For standard play, the exogenous retailer customer orders 4 units each round for the first 4 rounds, which means the retailer sees and fills orders of 4 units. The delay in order slips (and initial order slips of 4 units on the board at the time the game starts) also implies that all other players receive orders for 4 units in rounds 1-4. Thus, for each of the first 4 rounds, each player begins and ends with 12 units in inventory. This leads to team costs of $96 associated with training. ■

Theorem 3: Obtaining the minimum team score in standard game play requires the factory to drain the existing inventory, then produce exactly the number of units that will make inventory match expected customer orders.

Since the goal of the game is to minimize total team cost, and the other players act as a pass-through, this should be self-evident. ■

As shown in the system diagram in Figure 2, after the inventory drops to 0, the inputs to the inventory should equal the outputs, which occurs when factory orders exactly match what the customer will order in 8 rounds. The factory must use the available information when forecasting demand. Specifically, after the training rounds customer demand may change (e.g., for standard play, the customer orders increase according to a discrete time step function to 8 units per round in round 5). We can model the basic deck of customer demand for game play as:

$$O[r] = (k - O_0)u[r - r_0] + O_0$$

where k is a positive integer and $u[x]$ denotes the discrete-time unit step function with $u[r] = 1$ for $x \geq 0$ and $u[r] = 0$ for $x < 0$ (for the standard deck: $k = 8$, $r_0 = 5$, and $O_0 = 4$). We focus on decisions for factory ordering, starting at the time of placing orders autonomously at the end of round 4. The factory player controls all of the inventory in the system according to factory orders, $f[r]$. Based on the cost function and the objective of cost minimization the factory player should seek to maintain an inventory for the retailer such that $R[r] = 0$. Thus, at the time of placing a first production order in round 4, the factory should logically order (i.e., start production on) 0 units, $f[4] = 0$, because the system currently contains significant excess inventory (i.e., 12 units in inventory for each player). At this point in time, the factory should expect to continue to order 0 for some time period to reduce the inventory in the system to adjust for apparent low customer demand to date and based on an *a priori* assumption of constant customer demand (as occurred in the training rounds). However, customer demand (i.e., $O[r]$) remains uncertain from a decision analytic perspective, and consequently additional information about customer demand remains valuable.

With excess inventory in the system at the start of the game, the factory can infer customer orders by monitoring the retailer inventory (so long as the retailer possesses sufficient inventory to fill customer orders, see Theorem 2). With standard game play, the factory can observe the change in customer demand as it occurs in round 5 for $k < 17$ (i.e., $k > 16$ drops the retailer to 0 in round 5, and thus the factory would observe this as $k \geq 16$).

Due to delays in the system, it takes 8 rounds for new factory production to affect the retailer inventory. When determining production orders for the factory, $f[r > 3]$, the factory must consider all of the inventory in the system and seek to minimize the cost of system inventory based on the following recurrence relation:

$$R[r] = R[r - 1] + O[r] + S[r] + f[r - 8]$$

Where $R[-1] = 12$ represents the initial inventory in the retailer, and $S[r]$ represents the inventory in the system at the time of initialization that will flow into to the retailer after training, which equals: 4, 16, 4, 16, 4, 16, 4, 4 (i.e., the stream of units in the system assuming the other players optimally pass the inventory in the system up to the retailer) and is only defined for rounds 1 through 8 (i.e., zero after round 8). The factory order at the end of round 4 will affect the inventory of the retailer when all of the stock defined by $S[r]$ reaches the retailer, at which time $80 - 8k$ stock would exists in the system if the factory did not add any additional inventory (i.e., ordered production of 0 units). Thus, assuming that any change in inventory will remain constant (stationary) after the change, the factory player can calculate the expected inventory or backlog of the system based on the inferred value of k from monitoring changes in the retailer inventory. After observing the value of k change in round 5, the factory player can compute the appropriate amount of production to order assuming a stream of subsequent customer orders of k units per round. Thus, the strategy for the factory order depends on the value of k, which the factory should observe when $r = r_0$ (round 5):

- For $8 < k \leq 16$, order $10k - 80$ units in round 5 and order k units for subsequent rounds ($r > 5$). If the retailer filled an order of 16, the customer order may exceed 16, but this remains unobserved and unobservable.
- For $2 < k \leq 8$, the factory needs to wait until the supply chain contains $8k$ inventory before starting to send k units per round. For values of k that do not divide evenly into 80, the factory must add some extra inventory to the system the turn before the supply chain drops to $8k$ inventory so that the supply chain reaches exactly $8k$ inventory on the next turn.

The factory should continue to monitor inventory levels in the retailer (and thus infer customer orders filled each round). This allows the factory to adjust and account for any signals of non-constant demand. Given uncertainty and the potential for customer orders to change at any time, the factory could decide to maintain excess inventory at the retailer position (i.e., one or more extra units) such that depletion of this stock would provide a signal of an increase in customer orders, although this strategy incurs some additional inventory costs.

Theorem 3: With step-demand orders, perfect information only facilitates a lower team score if the initial conditions do not imply excess inventory required for k.

With perfect information (i.e., prior knowledge of k), the factory can improve its strategy for $k > 8$, because instead of ordering 0 units of production at the end of round 4, the factory can send stock to the retailer at the first opportunity to place an autonomous order, and thus reduce the backlog generated by the one-turn delay in the factory. If the factory player expects that $k > 8$ at the start of the game, the factory player can anticipate the demand by sending more inventory at turn 0, however, the training rounds do not motivate such *a priori* anticipation. ■

We can define the best strategy for any set of customer orders as the strategy that controls the system inventory at 0 for all rounds after depletion of the initial inventory associated with the game set up and training rounds. In some cases, a best strategy may not exist, because the factory cannot maintain the retailer at 0 inventory, such as the case of random customer demand. For very high customer demand (e.g., $k > 16$ in the step demand case), the factory will need to send a very large spike of inventory downstream and observe the retailer stock once it arrives to determine the amount of backlog filled and observe any pattern of demand. To generalize our observations, we define n as the number of rounds required to observe and predict a stationary or periodic (i.e., non-random) pattern of customer demand.

Theorem 4: For a non-random customer demand function characterized by n, the best strategy exists with an $r_0 \geq n+8$.

For a customer demand deck characterized by n, it takes at least n rounds until the factory player can observe the customer order pattern. After the factory player understands the pattern, the factory player can construct a strategy that controls retailer inventory at 0 (e.g., for a two-step function as used in the beer game expert deck). ■

The factory order strategy presented for the step demand deck holds for a very narrow demand function with $n = 1$. However, we can generalize the factory strategy along the same lines for a deck defined with a transient, $t[r]$, that settles to a periodic function, $P[r]$:

$$O[r] = t[r] + P[r]u[r - r_0]$$

In this case, $t[r] = 0$ for all $r \geq r_0$. The strategy for the factory relies on the ability of the factory to observe $P[r]$, and to observe the total stock depleted by the transient. In this case, the factory must supply enough inventory to handle the transient demand, and then quickly phase-lock to $P[r]$ with a fixed delay of eight rounds. In order to do this, the player at the factory should initially supply enough inventory for the maximum value of $P[r]$, so that the retailer does not go into backlog, because the cost function favors holding inventory over backlog and excess inventory remains observable, while backlog information remains censored.

Because of the generality of functions of this form, we cannot identify a closed-form strategy for a deck in which a periodic function follows a transient. However, the idea of the strategy remains the same: hold enough inventory (all at the retailer) to account for the transient component, and then send exactly the customer demand when this becomes known through observation of customer orders filled. Therefore, the factory player operates in two stages: first, the player assumes demand with large n during the transient while observing the customer demand, and second the player switches to controlling inventory at 0 after identifying the best strategy. During the second phase, the factory player can also maintain excess inventory in the supply chain to manage unexpected increases in demand. Because this is a linear system, any change in demand can be controlled by treating the disturbance as a second order function.

III. SIMULATION AND FIELD APPLICATION

We apply the strategies and equations developed in the prior section to characterize the factory orders that minimize the team score for our strategy for all customer order streams that go from a constant of 4 units per round for the initial set up and during the training rounds to a constant integer value (k) between 3 and 16 units per round for the remaining rounds. We estimate the minimum team scores for each k assuming 36 rounds total (including the training rounds) to provide a consistent comparison with prior results (although using our approach cumulative team scores for additional rounds of 0 mean that these rounds do not change the overall score). We constructed a discrete event simulation of the game in Excel that allowed us to verify the results from the equations above. To provide additional context, we also characterize the overall team costs that result from strictly passing the information about the customer orders up the supply chain [14]. We also consider potential strategies that would allow players, specifically the factory, to observe the

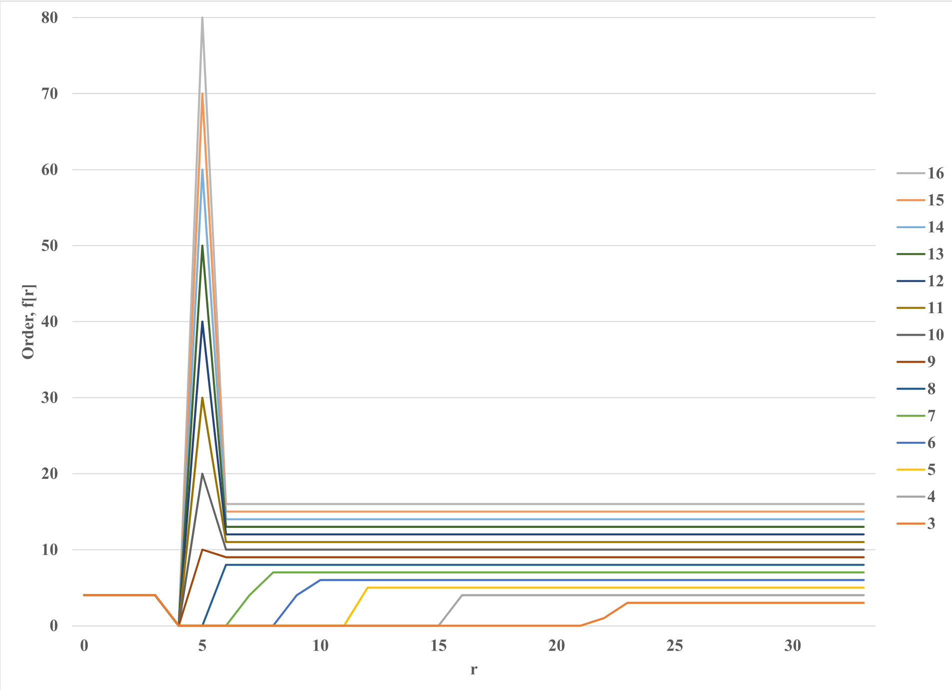

FIGURE 3. Optimal factory orders, *f [r]*, by round *r* for different values of customer demand *k*, including training for which the factory order 4 units for 3 rounds followed by an order of 0 units in round 4.

occurrence of a late change in the stream of retail customer orders following depletion of the inventory in the system. Uncertainty about actual customer orders may motivate the factory to send a pulse of inventory to temporarily increase retailer inventory or to leave one unit in excess inventory in the retailer position that would signal a change if the retailer used this small reserve. We further consider overall team costs if the factory benefits from perfect information about customer orders.

We field tested this strategy during the world beer game championship held at the 25th Annual International System Dynamics Society Conference in Cambridge, MA in July 2007. At the beginning of the competition, we met the other members who joined the team. We assumed the roles of the factory and the retailer, and we asked the other team members to play the game assuming that they should seek to minimize the cost of the inventory of the upstream player and assume responsibility for achieving and maintaining zero inventory for the distributor and factory, respectively.

IV. RESULTS

Figure 3 shows the stream of optimal factory orders, including the training rounds, and assuming use of the available imperfect information that minimizes the total team score for values of k between 3 and 16. As shown, for $k < 9$, the factory orders 0 when $r = 5$ and for various lengths of time (i.e., for 1 additional round for $k = 8$ up to 18 additional rounds for $k = 3$). For $k \geq 8$, Figure 3 shows the constant orders of k beginning at $r = 6$.

Figure 4a shows the minimum team total scores that correspond to the factory orders in Figure 3. We find application of this optimal strategy for the standard game (i.e., $k = 8$) leads to a total team score of \$148, which represents a considerable improvement over previously published benchmarks. We estimate the lowest possible overall team costs given the initial set up of the game of \$133, which occurs with a $k = 9$.

The squares in Figure 4b show the lower overall team costs associated with perfect information about customer orders for the factory compared to the circles that show the optimal strategy with the imperfect information available just from observing the board (i.e., the results shown in Figure 4a). The results show the same values for $k < 9$. This occurs because perfect information makes no difference in the minimum overall team score if the supply chain starts out with significant excess inventory. This implies an expected value of perfect information of \$0 for the standard game. For $k > 8$, the ability to order optimally from the time of the first autonomous order (i.e., when $r = 4$) instead of ordering 0 at that point based on an *a priori* assumption of expected excess inventory slightly reduces the cost, with the difference between the curves increasing with higher values of k. Even for $k = 16$, the expected value of perfect information is only \$64. Table 1 summarizes the team costs with available and perfect information, the expected value of perfect information, and the team costs of the strategy of passing customer order information up the supply chain, which leads to significant backlog for all players and consequently to much higher costs. For the beer game expert deck, our strategy gives an

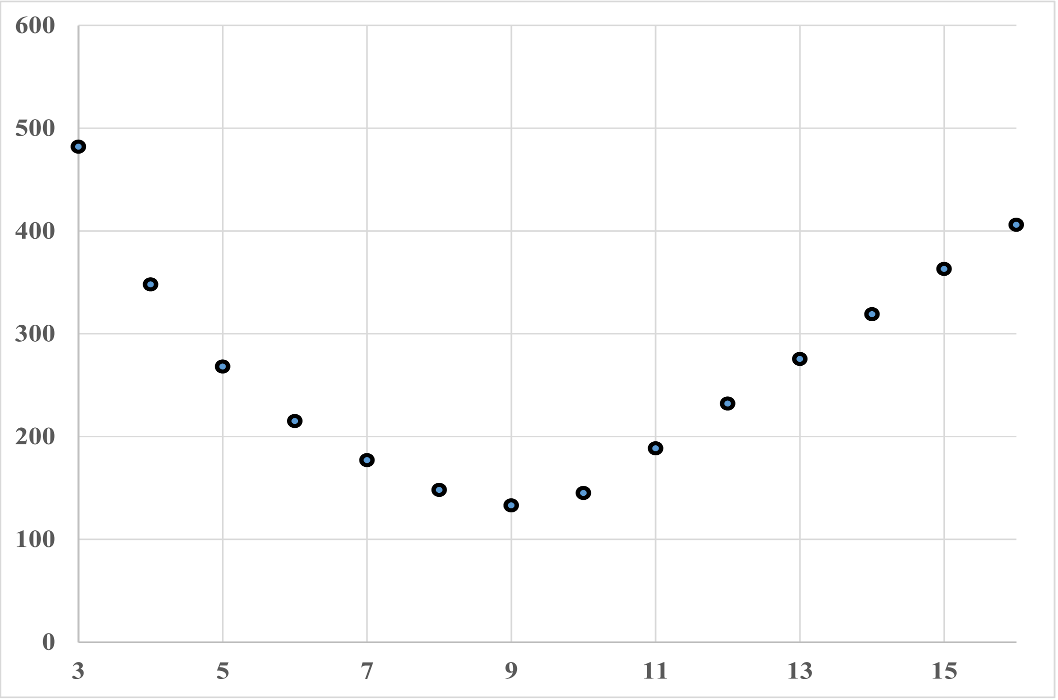

(a)

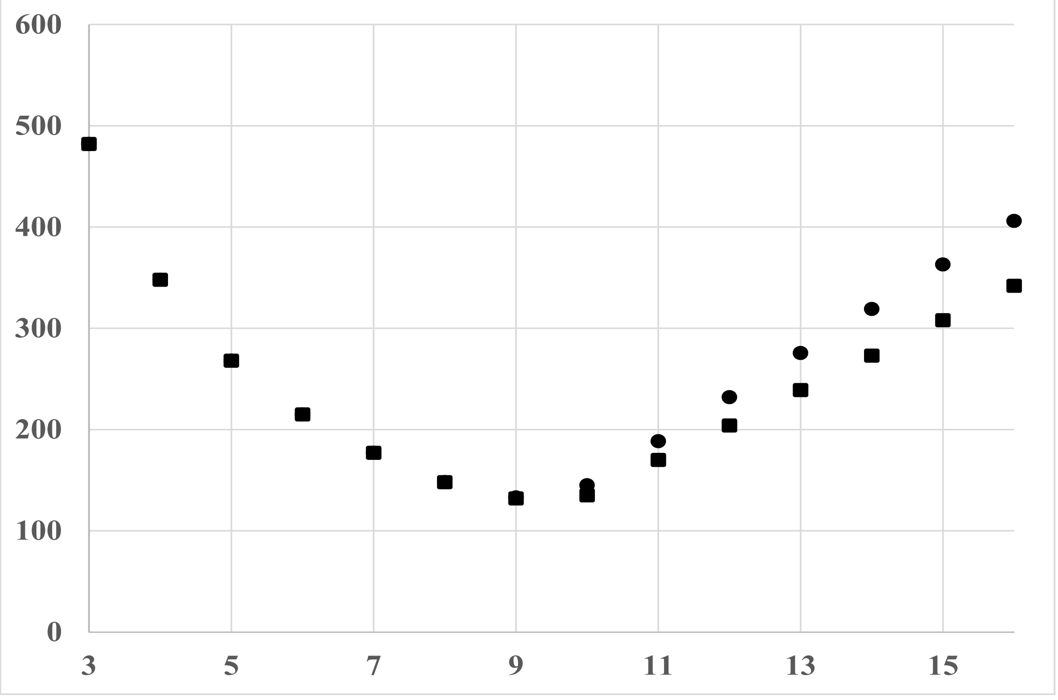

(b)

FIGURE 4. (a) Minimum overall team score including costs of training as a function of *k* for our optimal strategy with available imperfect information. (b) Minimum overall team score including training as a function of *k* for our optimal strategy with available imperfect information (circles) and with perfect information (squares).

overall team score of \$202 using available information and a score of \$178 with perfect information, which implies an expected value of perfect information of \$24.

During the field application of our strategy, the customer orders used $k = 4$, implying a lowest possible minimum team cost of \$348. We won the competition using our strategy,

TABLE 1. Minimum overall team cost including training as a function of *k* for our optimal strategy with available imperfect information, with perfect information, and for customer orders passed up the chain, and the expected value of perfect information.

k	Cost with available information ($)	Cost with perfect information ($)	Expected value of perfect information ($)	Cost of passing information up ($)
3	482	482	0	1073
4	348	348	0	864
5	268	268	0	656
6	215	215	0	447
7	177	177	0	239
8	148	148	0	792
9	133	132	1	1656
10	145	135	10	2520
11	188.5	170	18.5	3408
12	232	204	28	4296
13	275.5	239	36.5	5184
14	319	273	46	6072
15	362.5	308	54.5	6960
16	406	342	64	7848

with a score of $372. During game play the retailer, wholesaler, and distributor implemented the optimal strategy of maintaining 0 inventory of the players to their right without deviation and obtaining the expected costs. The factory followed the strategy of ordering 0 units to deplete the supply chain prior to ordering k units. Given uncertainty about the customer orders, for the actual game after several rounds with 0 inventory in the retailer, the factory sent a surge of inventory (e.g., 16 units one round followed by 0 units for 3 subsequent rounds) at one time to ensure that the retailer remained out of backlog. This approach provided confidence that customer orders had remained steady at $k = 4$, which added a small level of additional costs to the total team costs, but still allowed the team to win by a considerable margin.

V. DISCUSSION

Despite decades of insights derived from the beer distribution game, our analysis and experience demonstrate that simple strategies for each player and better use of the available information lead to significantly lower expected costs than identified as benchmarks in prior studies. Teams can easily implement this strategy in the context of game play without any additional information, which we demonstrated in a real situation. For this relatively simple supply chain, understanding the system, decision options, and information available from observation reduces the complexity significantly compared to attempts by players to manage their own inventory despite their limited control. In addition, the strategy requires that only the factory player needs to make *a priori* assumptions about the retailer customer orders, and that the factory player alone needs to carefully observe and control the total inventory in the supply chain.

While this approach may appear unrealistic in real supply chains, consistent with the unrealistic aspects of this stylized game [8], [15], some aspects warrant consideration. First, the importance of good information about customer orders in managing supply chains continue to drive actual manager behavior, and retailers can work directly with factories in some cases to either help the factories produce appropriate amounts of inventory to meet consumer demand or by setting prices that can potentially influence consumer demand in a way that matches factory production. Second, starting in the 1970s, large retailers recognized a significant advantage from buying directly from manufacturers and aggressively managing information about their customer orders. Ironically, Wal-Mart apparently began its efforts in this area out of necessity, not from a conscious decision, but very early on the senior management appreciated the "brilliance" of the approach [3]. By managing all aspects of distribution within their supply chains retailers can move inventory to sell it faster, which allows them to use the revenue as interest-generating capital for the time period between the sale of the product and payment of the bill for the inventory purchased from the factory. Third, within the construct of the simulation game, players can operate as a team and avoid the bullwhip effect, but this differs from acting to attempt to optimize their own inventories, which they do not actually control. The focus on operations can lead to deep change and significant strategic advantage for companies [29], but it requires truly understanding the system. Our analysis underscores the importance of understanding the impact of individual and collective decisions in supply chains (and complex systems more broadly) and the extent to which participants face incentives that will lead them to make better or worse decisions for the parts that they control. Fourth, while we cannot expect all players to behave rationally [11], [15], the optimal strategy creates the right incentives. In real supply chains, although individual players may receive rewards only for the costs they incur and they may make decisions that negatively impact others in the supply chain, from a societal or welfare perspective the system costs and total supply chain costs matter.

The beer distribution game provides an interesting case in which perfect information may not improve decisions significantly or at all (i.e., for $k < 9$). Instead, optimal performance requires using the available imperfect information well and understanding the system. Following up on seminal empirical work [11], more recent empirical studies of the beer game with stationary demand demonstrated that teams given perfect information still made decisions that induced the bullwhip effect, although somewhat dampened for some experimental treatments [30], [31]. This suggests that perfect information about demand is not the real issue, and that poor performance results from failure to understand the system and individual and organizational behavior.

This analysis of a relatively simple game played over a period of 50 years, for which prior publications all suggested worse ''best'' scores than the optimal scores we identified using our integrated analytical approach, reveals new insights and contributes to the body of evidence that highlights the opportunity for better performance by those who understand the system, information available, and player/stakeholder incentives. For example, in the last decade, individuals and organizations that understand the dynamics of the financial trading system identified optimal strategies for performance that significantly changed the incentives for all stakeholders in the market, [32] and sports teams use their understanding of the system and players to field better teams and win games [33]. The value of asking the question about what one would and should do with perfect information may prove more valuable with respect to motivating understanding of the system and lead to better decisions and performance than actually receiving the information itself.

ACKNOWLEDGMENTS

The authors thank Professor John Sterman for introducing us to the beer game, and Drs. John Lyneis and Brandon Owens for playing the wholesaler and distributor roles during the field application. They dedicate this paper to Professor Jay Forrester, whose vision and contributions continue to inspire students, teachers, and managers around the world.

REFERENCES

[1] J. Sterman, *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Boston, MA, USA: McGraw-Hill, 2001.
[2] K. M. Thompson, R. J. D. Tebbens, M. A. Pallansch, S. G. F. Wassilak, and S. L. Cochi, ''Polio eradicators use integrated analytical models to make better decisions,'' *Interfaces*, vol. 45, no. 1, pp. 5–25, Feb. 2015.
[3] D. Glass and H. Gilman, ''The most underrated CEO ever: The legendary Sam Walton got the credit, but it was David Glass who turned Wal-Mart into the world's largest company. In an exclusive interview, Glass talks about the early days, the boom years, and how China will keep the mega-merchant going strong,'' *Fortune Mag.*, vol. 149, no. 7, pp. 242–245, Apr. 2004.
[4] G. Stalk, P. Evans, and L. E. Sgulman, ''Competing on capabilities: The new rules of corporate strategy,'' *Harvard Bus. Rev.*, vol. 63, pp. 57–68, Mar./Apr. 1992.
[5] A. Latour, ''Trial by fire: A blaze in Albuquerque sets off major crisis for cell-phone giants. Nokia handles supply shock with aplomb as Ericsson of Sweden gets burned,'' *Wall Street J.*, p. A1, Jan. 2001.
[6] B. Bakken, J. Gould, and D. Kim, ''Experimentation in learning organizations: A management flight simulator approach,'' *Eur. J. Oper. Res.*, vol. 59, no. 1, pp. 167–182, May 1992.
[7] D. C. Lane, ''On a resurgence of management simulations and games,'' *J. Oper. Res. Soc.*, vol. 46, no. 5, pp. 604–625, May 1995.
[8] E. G. Anderson, Jr., and D. J. Morrice, ''A simulation game for teaching service-oriented supply chain management: Does information sharing help managers with service capacity decisions?'' *Prod. Oper. Manage.*, vol. 9, no. 1, pp. 40–55, 2000.
[9] J. W. Forrester, *Industrial Dynamics*. Cambridge, MA, USA: MIT Press, 1961.
[10] International System Dynamics Society. *Production Distribution Game—'The Beer Game'*. [Online]. Available: http://www.systemdynamics.org/products/the-beer-game/, accessed Nov. 11, 2015.
[11] J. D. Sterman, ''Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment,'' *Manage. Sci.*, vol. 35, no. 3, pp. 321–339, Mar. 1989.
[12] R. Mason-Jones and D. R. Towill, ''Information enrichment: Designing the supply chain for competitive advantage,'' *Supply Chain Manage., Int. J.*, vol. 2, no. 4, pp. 137–148, 1997.
[13] F. Chen and R. Samroengraja, ''The stationary beer game,'' *Prod. Oper. Manage.*, vol. 9, no. 1, pp. 19–30, Mar. 2000.
[14] S. M. Hong-Minh, S. M. Disney, and M. M. Naim, ''The dynamics of emergency transhipment supply chains,'' *Int. J. Phys. Distrib. Logistics Manage.*, vol. 30, no. 9, pp. 788–816, 2000.
[15] J. D. Sterman and G. Dogan, '''I'm not hoarding, I'm just stocking up before the hoarders get here': Behavioral causes of phantom ordering in supply chains,'' *J. Oper. Manage.*, Jul. 2015. [Online]. Available: http://dx.doi.org/10.1016/j.jom.2015.07.002
[16] J. W. Pratt, H. Raiffa, and R. Schlaifer, *Introduction to Statistical Decision Theory*. Cambridge, MA, USA: MIT Press, 1965.
[17] R. A. Howard, ''Information value theory,'' *IEEE Trans. Syst. Sci. Cybern.*, vol. 2, no. 1, pp. 22–26, Feb. 1966.
[18] R. A. Howard, ''Value of information lotteries,'' *IEEE Trans. Syst. Sci. Cybern.*, vol. 3, no. 1, pp. 54–60, Jun. 1967.
[19] H. Raiffa, *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*. Reading, MA, USA: Addison-Wesley, 1968.
[20] F. Yokota and K. M. Thompson, ''Value of information literature analysis: A review of applications in health risk management,'' *Med. Decision Making*, vol. 24, no. 3, pp. 287–298, Jun. 2004.
[21] P. I. Frazier and W. B. Powell, ''Paradoxes in learning and the marginal value of information,'' *Decision Anal.*, vol. 7, no. 4, pp. 378–403, Oct. 2010.
[22] D. Bhattacharjya, J. Eidsvik, and T. Mukerji, ''The value of information in portfolio problems with dependent projects,'' *Decision Anal.*, vol. 10, no. 4, pp. 341–351, Dec. 2013.
[23] L. Steuten, G. van de Wetering, K. Groothuis-Oudshoorn, and V. Retèl, ''A systematic and critical review of the evolving methods and applications of value of information in academia and practice,'' *Pharmacoeconomics*, vol. 31, no. 1, pp. 25–48, Jan. 2013.
[24] J. M. Keisler, Z. A. Collier, E. Chu, N. Sinatra, and I. Linkov, ''Value of information analysis: The state of application,'' *Environ. Syst. Decisions*, vol. 34, no. 1, pp. 3–23, Mar. 2014.
[25] F. Chen, ''Information sharing and supply chain coordination,'' in *Handbooks in Operations Research and Management Science*, vol. 11, A. G. de Kok and S. C. Graves, Eds. Amsterdam, The Netherlands: Elsevier, 2003, pp. 341–421.
[26] J. Dejonckheere, S. M. Disney, M. R. Lambrecht, and D. R. Towill, ''The impact of information enrichment on the bullwhip effect in supply chains: A control engineering perspective,'' *Eur. J. Oper. Res.*, vol. 153, no. 3, pp. 727–750, Mar. 2004.
[27] Y. Ouyang, ''The effect of information sharing on supply chain stability and the bullwhip effect,'' *Eur. J. Oper. Res.*, vol. 182, no. 3, pp. 1107–1121, Nov. 2007.
[28] H. L. Lee, V. Padmanabhan, and S. Whang, ''The bullwhip effect in supply chains,'' *Sloan Manage. Rev.*, vol. 38, no. 3, pp. 93–102, 1997.
[29] M. Hammer, ''Deep change,'' *Harvard Bus. Rev.*, vol. 82, no. 4, pp. 84–93, 2004.
[30] R. Croson and K. Donohue, ''Behavioral causes of the bullwhip effect and the observed value of inventory information,'' *Manage. Sci.*, vol. 52, no. 3, pp. 323–336, Mar. 2006.
[31] R. T. A. Croson, K. L. Donohue, E. Katok, and J. Sterman, ''Order stability in supply chains: Coordination risk and the role of coordination stock,'' *Prod. Oper. Manage.*, vol. 23, no. 2, pp. 176–196, 2014.

[32] M. Lewis, *Flash Boys: A Wall Street Revolt*. New York, NY, USA: Norton, 2014.
[33] M. Lewis, *Moneyball: The Art of Winning an Unfair Game*. New York, NY, USA: Norton, 2003.

KIMBERLY M. THOMPSON received the B.S. and M.S. degrees from the Massachusetts Institute of Technology (MIT), and the D.Sc. degree from the Harvard School of Public Health. She founded and serves as the President of Kid Risk, Inc., a non-profit organization that uses analytics to improve children's lives. She is currently a Professor of Preventive Medicine and Global Health with the College of Medicine, University of Central Florida. Prior to moving to FL, she served as a Faculty Member with the Harvard School of Public Health for over a decade and spent two years as a Visiting Faculty Member with the MIT System Dynamics Group.

NIMA D. BADIZADEGAN is currently pursuing the B.S. degree in electrical engineering with the California Institute of Technology. His research interests include digital and mixed signal circuits. He is a member of the IEEE Solid State Circuits Society.