The Mexico-China Sourcing Game: Teaching Global Dual Sourcing[1]

Gad Allon & Jan A. Van Mieghem
Kellogg School of Management
Northwestern University
2001 Sheridan Road, Evanston, IL, 60208
g-allon@kellogg.northwestern.edu, vanmieghem@kellogg.northwestern.edu

Jul 20, 2009; Revised Nov 9, 2009; Jan 18, 2010; Jan 25 2010; Feb 10, 2010

We describe a three-hour class on global dual sourcing built around a game that demonstrates the challenges in making operational decisions and transfers recent academic insights to the classroom. Student teams manage a firm with access to a responsive but expensive supply source (Mexico) and a cheap but remote source (China). Each team must determine a sourcing strategy to satisfy random demand that is revealed throughout the game. In each period, teams place orders to both sources and manage two assets: inventory and their bank account. The goal is to maximize each team's value (final bank balance). During the debriefings, we analyze the policies used by different teams along both financial and operational metrics, present the optimal strategy, and summarize the experiential learning points.

Keywords: dual sourcing, strategic sourcing, experiential learning, inventory management, total landed cost, simulation game.

1. Introduction

In many retail settings, and business in general, firms can source goods from multiple suppliers. We investigate a typical choice between two suppliers: One is low cost but has long lead times while the other provides quick response but at higher cost. This classical problem is faced by many companies and most students understand its relevance. However, effective management of dual sourcing is surprisingly challenging and conveying its complexity through traditional means such as a case study or a lecture is difficult: Solving a dual sourcing case requires analytic tools that are not readily available to students or instructors. A lecture on dual sourcing runs the risk that students may not appreciate the added level of complexity in day-to-day, as well as strategic, dual sourcing relative to single sourcing.

We have developed a sourcing game as an experiential learning tool that addresses the above shortcomings. In this game, students play the role of sourcing managers who must make strategic allocation decisions as well as place day-to-day orders to two suppliers. During this process, they experience the operational, financial and service related consequences of their decisions.

[1] Submitted to *INFORMS Transactions on Education*'s special issue on Teaching Service and Retail Operations Management

The sourcing game is designed with four pedagogical objectives in mind:

1. To develop intuition on possible sourcing strategies and to appreciate simple, robust policies to guide sourcing decisions.
2. To highlight the role of working capital estimation in the concept of total landed cost.
3. To appreciate the added complexity of managing a supply portfolio over single sourcing.
4. To understand the futility of guessing and over-reacting to demand.

The game has been used with MBA and executive education students and although some knowledge of basic inventory models is useful, it is not necessary.
In the remainder, we will describe sequentially the three-hour class we have taught in three parts: Part A introduces global sourcing and related theory; Part B describes the game; and Part C summarizes the debriefing phase as well as a discussion of recent academic research. We conclude this paper by summarizing student reactions and possible extensions.

2. Class Part A: Introduction to Global Sourcing

This section describes the first part of the class that introduces global sourcing and related theory (about 30 to 45 minutes of class time).

2.1. The Motivating Case

As preparation for the class, the students are asked to read the Mexico-China mini-case in Van Mieghem (2008, pp. 230-232). The case describes a $10 billion high-tech U.S. manufacturer of wireless transmission components. The company was at a crossroads regarding its global network; the case focuses on its two assembly plants, one in China and the other in Mexico. While the Chinese facility enjoyed lower costs, ocean transportation made its order lead-times 5 to 10 times longer than those from Mexico. With highly uncertain product demand--coefficients of variations of monthly demand for some products were as high as 1.25--sole sourcing was unattractive: Mexico was too expensive and China too unresponsive.
The firm had to decide how it could best utilize these two sources. At the strategic level, this amounts to properly allocating product demand to each source. Strategic allocation refers to the expected cumulative demand allocated to each source over the planning horizon. At the tactical level, the firm had to choose a dynamic ordering policy that implements that strategic allocation at lowest cost. In practice, specifying strategic allocations and ordering policies are key tasks of any sourcing strategy--be it global or domestic--because it affects costs and supplier management.

2.2. Discussion on practice:

To connect the game to real sourcing practices, and if students have prior experience with similar settings, we start with an open discussion covering two questions:

1. *What policies have you used in practice*? This discussion typically reveals some of the following practices:
    a. Allocate products to plants based solely upon historic allocations (e.g., new products within a family are allocated to plants in the exact same fashion as existing products in that family regardless of supply and demand characteristics).
    b. Allocate products to plants based upon internal politics (e.g., "pet" production facilities get higher volume, lower complexity products to keep utilization high and cost structure low).
    c. Allocate products to plants based upon basic understanding of costs (e.g., China's cost base is lower – thus source full requirements offshore).
    d. Product allocations follow a simple primary-secondary allocation (e.g., if primary location cannot fulfill demand – requirements cascade to the secondary location).
    e. Product-to-plant allocation decisions are rarely re-visited or adjusted (e.g., allocation decision is a "life sentence").
2. *What data/metrics are used to make sourcing decisions*? This discussion follows the practices above and highlights the typical cost, leadtimes, and service risk metrics. The instructor than can ask the natural question: "How can we combine these various components into a single metric?" This brings us to the concept of total landed cost.

2.3. The Concept of Total Landed Cost (TLC)

To address the motivating case, the natural first step is to compare the two extreme solutions: single-sourcing from Mexico versus single-sourcing from China. This requires accounting for the difference in cost-of-good-sold, shipping costs, duties, and working-capital requirements[2]. For that purpose we review the concept of total landed cost (TLC), which represents the end-to-end cost to transform inputs at the source to outputs at destination. The TLC captures not only the traditional cost of goods sold, but also accounts for supply chain costs such as transportation, customs, duties, taxes, as well as required working capital carrying costs to achieve desired service levels and protect against supply and demand risks; see Fig 1.

[2] It should be noted that working capital includes the pipeline inventory as well as the safety stock. The pipeline inventory depends on the transformation process, which takes place between the moment of ordering and the moment of receipt of the order. If part of the lead-time is caused by inflexibility in production due to set-up times, then there is no per-se investment in material during the lead-time period. Thus, in order to accurately assess the working capital requirements one needs to know the precise timing of material inflow to build a sequence of cash flows over time.

We will refer to all but working capital cost components as the "sourcing cost." Computing the sourcing cost is tedious yet straightforward. In contrast, working capital is driven by pipeline and safety inventory, which depend on lead-times, volatility and service levels.

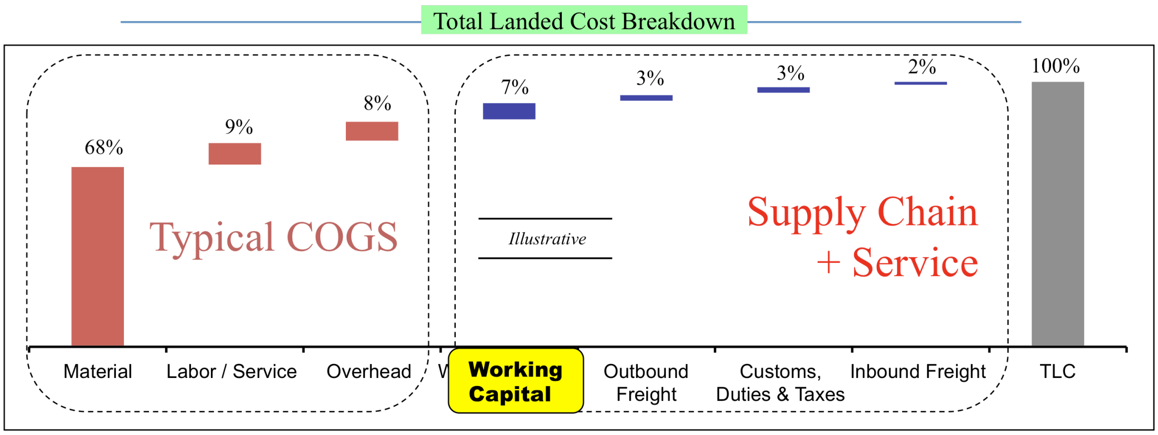

Figure 1: Total Landed Cost (TLC) is the total cost incurred from source to destination.

2.4. Computing TLC for single sourcing

Working capital and inventory requirements are easily estimated for single sourcing using readily available standard inventory formulae. Before moving to dual sourcing and to the game, we remind the students of the concept of safety stock and compute the total landed cost for each of the locations separately. This computation reveals that single sourcing from China appears to be the favorite solution.

Asking the students' reaction to this analysis generates a discussion of the merits of sourcing from Mexico and of the expected benefits of dual sourcing in a quest to combine the strengths of both sources. This also leads to a discussion on the risk of using a single source. We usually discuss two types of supply risks: stochastic lead times and supply disruptions (such as natural disasters strikes.) This naturally leads to the question of *how* to design an effective dual sourcing policy, which is the transition point to the game.

3. Class Part B: The Mexico China Sourcing Game

This section covers the main part of the class by describing the different stages of the game (about 1 to 1 ½ hour of class time).

3.1. Game setup (mostly done in advance of class)

The idea behind the game is for students to act as sourcing managers of a firm making periodic ordering decisions for a new product. Demand is highly uncertain and a probabilistic forecast is provided.

Setting: A class with as few as 5 or as many as 60 students divided into 5 to 10 groups.

Props: Each student group needs (at least) one laptop computer. Each group receives a log-in code for the web-site hosting the game[3]. The instructor needs a laptop that is connected to the Internet and to a classroom screen projector.

Setup: Similar to the Mexico-China mini-case, each team represents an identical company that is introducing a new product that can be sourced from China or Mexico. Each team manages two assets: cash and inventory.

The demand distribution for the product is known and shown to all groups in advance. (The specific distribution that we have used is Gamma with mean 10 and standard deviation 15, reflecting a highly volatile demand.) The actual demand realization is identical for all groups and projected onto the classroom screen dynamically over time; see Fig 2.

Each team can place an order to Mexico and an order to China in each period; sources differ in cost and in leadtimes. Specifically, orders placed to Mexico have a sourcing cost of $8,000 per unit and arrive at the beginning of the subsequent period. Orders placed to China have a sourcing cost of $7,000 per unit but arrive only after 4 time periods. The product is launched in the 4th period, allowing the groups to build pipeline inventory from both sources before sales start. The product sells for $10,000 per unit. The starting version of the game assumes that any unit demanded but not available on-hand is backlogged for a cost of $20 thousand per unit and each team plays in a separate yet identical market.

All teams start with zero money in the bank and can borrow to finance inventory. Any bank balance (both negative and positive) incurs a 1% interest rate per period. All team bank balances are projected on the classroom screen continuously (which allows teams to compare their position to other teams and creates a sense of competition and excitement). Any leftover inventory at the end of the game is liquidated at zero value (although the instructor can easily change the salvage value).

The objective for the teams is to maximize their firm value, which is their bank balance, at the end of the game. The termination time of the game is determined by the instructor but is not told to students in advance to prevent end-of-horizon effects).

[3] For access to the game, please contact the authors.

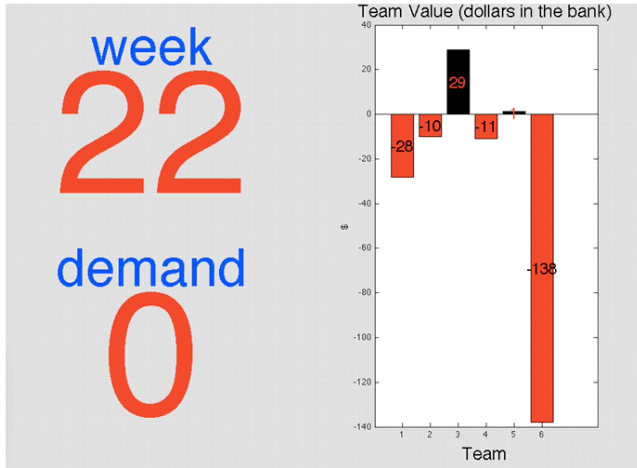

Figure 2: During the game, the time as well as corresponding demand and bank balances are projected on the classroom screen.

3.2. Kick-off and Brainstorming (between 15 to 30min of class time).

The goal of this stage is to allow each group to develop a sourcing strategy. Each group is given an Excel spreadsheet that allows students to simulate and test their strategies. Specifically, the spreadsheet shows how the pipeline inventory, sales, and bank balance evolve given user-entered orders and demand.
Groups are asked to think through how they would react (in terms of placing orders) given a new demand realization and a certain inventory status. To steer their thinking, each group is also asked to estimate the fraction of orders they will place to China and to Mexico and hand this in to the instructor.

3.3. Playing the dual sourcing game (between 45 to 60min of class time)

The game starts with a four period "pre-launch stage" during which students can fill their pipeline before sales start. Demand and sales start following the pre-launch stage. During each period, all groups are informed simultaneously of the same period demand as shown in Fig. 2. (Showing all students the same demand stream creates a positive class atmosphere, manifested by students shouting "YES!" or "Oh Jeez" etc., where all students together experience the game.)

In each period, the typical inventory actions take place:

(i) Teams observe the period demand (which is zero during pre-launch);
(ii) Sales are determined automatically as the minimum of period demand and on-hand inventory (see Fig. 3 for a screenshot of the team's web browser interface);
(iii) Previous orders are received automatically;
(iv) Teams place orders to the two sources;
(v) Each team's pipeline inventory status and its bank balance are updated

automatically for the next period.

The instructor terminates the game at a time of his or her choosing. (We have typically played 30 to 35 periods, for about 50-60 minutes.) It is critical not to inform students of the termination period in advance to prevent end-of-horizon effects.

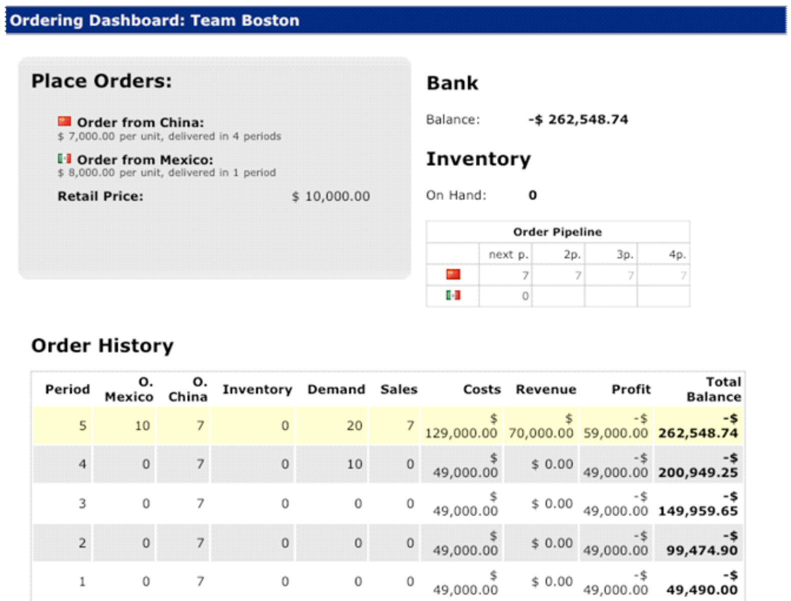

Period	O. Mexico	O. China	Inventory	Demand	Sales	Costs	Revenue	Profit	Total Balance
5	10	7	0	20	7	$ 129,000.00	$ 70,000.00	-$ 59,000.00	-$ 262,548.74
4	0	7	0	10	0	$ 49,000.00	$ 0.00	-$ 49,000.00	-$ 200,949.25
3	0	7	0	0	0	$ 49,000.00	$ 0.00	-$ 49,000.00	-$ 149,959.65
2	0	7	0	0	0	$ 49,000.00	$ 0.00	-$ 49,000.00	-$ 99,474.90
1	0	7	0	0	0	$ 49,000.00	$ 0.00	-$ 49,000.00	-$ 49,490.00

Figure 3: Screenshot of a team's web browser interface where students can place orders and track their performance.

4. Class Part C: Debrief and Connection to Recent Research

This section describes the last part of the class during which we debrief the game insights, discuss recent academic research, and summarize take-aways (about 30min of class time).

4.1. Debriefing the game

After ending the game, the instructor asks different groups (including the "winning" team, the "runner-up," and a low performer) for their sourcing strategy. Teams are encouraged to share their thoughts and experiences on the benefits of their strategy and what they would do differently next time.
During that discussion, every argument is validated against metrics compiled during the game. The instructor shows a "dashboard" (Fig. 4) that summarizes the performance of each team along financial and operational metrics:

(i) Financial metrics include the value of each team's firm (bank values at the end of the game), costs (total number of units sourced), and revenues (the cumulative number of units sold), and

(ii) Operational metrics include each team's actual strategic allocation (fraction of total units sourced from each location), average on-hand inventory, and service level measured by the fill rate (total number of units sold divided by the total demand).

We are also able to show the actual orders placed over time from each source for each group.

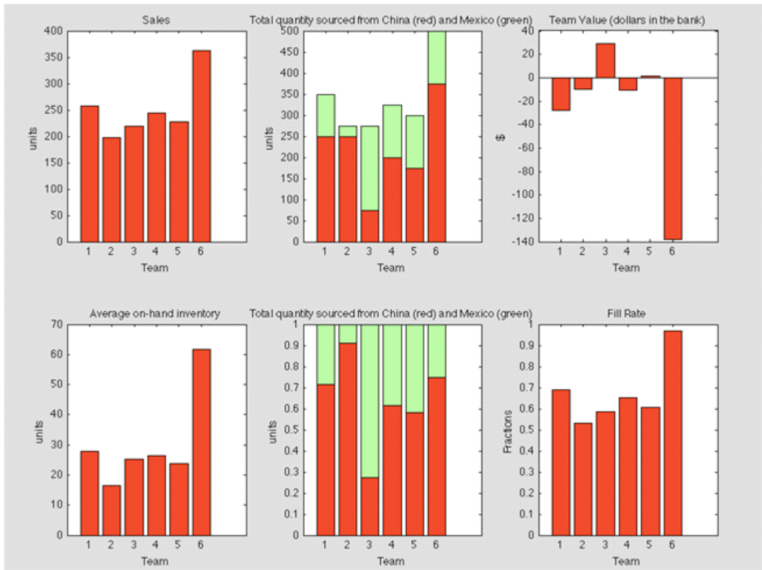

Figure 4: A dashboard that summarizes financial (top row) and operational (bottom row) metrics of all teams is projected on the classroom screen during the debriefing of the game.

The discussion leads to the observation that different teams may perform equally well financially while employing different strategies. The groups usually differ with respect to average on-hand inventory as well as the fraction of units sourced to China vs. Mexico. The common thread for all "successful" groups is the ability to execute the strategy and not change it in the face of lower (or higher) than expected demand or inventories. This naturally leads to the following questions: What are the features of a good dual sourcing policy? What guidelines can academic theory provide? Answering these questions brings us to the next topic.

4.2. Recent Academic Research: TBS and dual index policies

The objectives of the game include answering the following two questions: (a) what features do effective dual sourcing policies share? And (b) how can we determine near-optimal policies and average sourcing allocation?

Effective dual sourcing policies take into account not only the on-hand inventory and the current demand, but also the entire pipeline status as well as the entire demand forecast. Research has shown that the structure of the optimal policy is very complex. Therefore, the dual sourcing literature has traditionally focused on determining sophisticated dynamic policies that approach optimal performance. Typically, these policies are characterized by one or two target inventory levels (base-stock levels) and keep track of one or two inventory positions (indices). For example**,** Veeraraghavan and Scheller-Wolf (2006) show that a dual-index policy with two target levels performs close to optimally, which represents state-of-the-art dual sourcing research.

Unfortunately, determining the target levels of sophisticated policies such as the dual index policy requires sophisticated computational work (optimization via simulation). In addition, under these policies, the associated strategic sourcing allocation can be determined only through simulation. For a strategic sourcing manager, it would be desirable to have a simple policy with simple guidelines that determine the strategic allocation during the planning phase without significantly compromising performance.

Recently, Allon and Van Mieghem (2009) have presented such a policy that is used in practice: the tailored base-surge (TBS) policy. Under this policy, a constant ("standing") order is placed to the low cost supplier in each period. The responsive source is used only to bring the total inventory (on-hand + pipeline) to a single target level. The TBS policy thus echoes a fundamental tenet in strategy: it aligns the ordering patterns with the core competencies of the suppliers (Fig. 5). The constant base allocation allows China to focus on cost efficiency while Mexico's quick response is utilized only dynamically to guarantee high service.

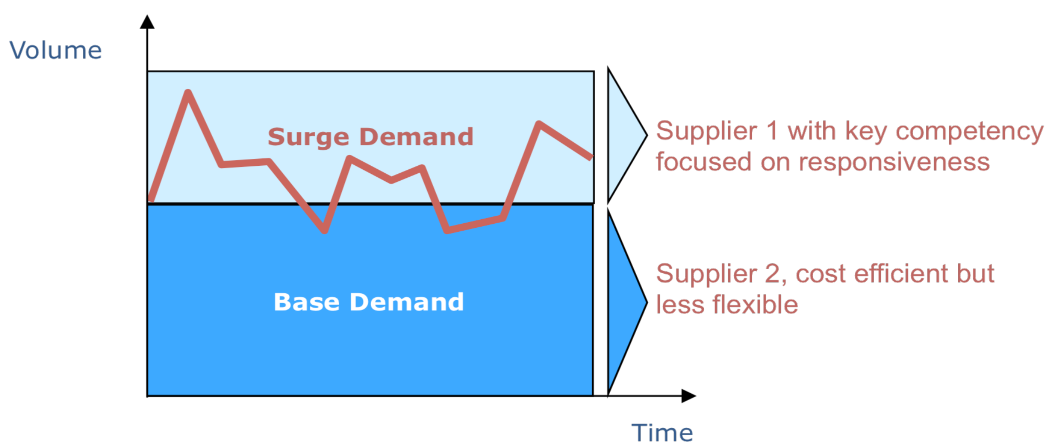

Figure 5: A Tailored Base-Surge Policy orders a constant fraction from the low cost source and orders from the responsive source to respond to surges.

Besides its simplicity, the TBS policy directly determines the strategic allocation, which equals the standing order divided by the average demand. Janssen and De Kok (1999) study a similar policy to the TBS and show using numerical study that a firm should allocate close to 90% of its demand to the cheaper source. Allon and Van Mieghem (2009) optimized the TBS policy analytically and present a simple square-root formula to specify the near-optimal strategic allocation:

$$\text{base allocation} \cong 1 - \sigma\sqrt{\frac{h}{2\Delta c\lambda}},$$

where *h* is the unit holding cost, Δc is the unit sourcing cost differential, λ is the demand rate, and σ is the supply-demand volatility.[4] Simple formulae also exist for the corresponding target inventory level and cost. When focusing only on demand volatility, the formula simplifies further to:

$$\text{base allocation} \cong 1 - \sqrt{\frac{h}{2\Delta c}}\text{COV}_D,$$

where COV_D is the coefficient of variation of the single-period demand distribution. The formula also provides insight by identifying the key drivers of dual sourcing and quantifying their interaction. Specifically, the allocation to the low cost supplier is high when:

1. The key monetary trade-off $\Delta c/h$ is large, which implies a large cost advantage or a small cost of capital (small commercial risk)
2. The expected demand rate λ is high
3. The supply-demand volatility σ is small. This arises with stable or high-volume products (in the maturity phase of the product life cycle) and requires stable (level) production in China. This factor also shows how the sourcing allocation should change as a product moves through its product life cycle.

Students are asked to compute the strategic allocation using the square-root formula for the parameters of the game, which is easy:

- *h* = about $7,000*1%/period = $70/period
- Δc = $1,000
- COV_D = 1.5

so that the near-optimal base allocation is approximately 1 – √(.07/2)*1.5 = 1 - .28 = .72 (As a home-work assignment, the instructor can ask for a similar computation for each of the different SKUs in the mini-case.).

In the class we also compare the performance of different classical policies for a finite horizon problem, similar to the game. Table 1 compares the profits of 4 strategies: single sourcing from Mexico, single sourcing from China, the tailored base-surge policy, and a dual base-stock policy. For each policy, we have optimized via simulation over the policy parameters. For single sourcing policies we computed the optimal base-stock policy. For the Tailored Base-Surge policy we used the formula given in class, and for the dual base-stock policy we used, again, simulation-based optimization to compute the optimal threshold levels. For each policy we present both the value a firm would have obtained using that specific policy and the specific demand sample path played in the game, as well as the expected value corresponding to the demand distribution.
We also compare these policies with the case in which the firm has perfect demand information. With perfect demand foresight, the optimal policy is simple: order the

---

[4] Specifically, σ^2 is the sum of the squared coefficients of variation of the interdemand times and the China supply times.

exact demand quantities four periods early from China. The comparison with perfect information highlights the dramatic cost of uncertainty.

Table 1

Strategy	Optimal Target Inventory Levels	Actual Value | game demand	Expected Value
Single Source Mexico	23	$557	$422
Single Source China	48	$423	$569
Tailored Base Surge	18, standing order to China = 5	$616	$547
Dual Base Stock	10, 45	$514	$586
Perfect information + Single source China	Order exact demand four periods earlier from China	$2,321	$1,878

As one may expect, the Dual Base Stock policy outperforms the other policies, in expectation. However, on the specific sample path played in the game, the TBS generates a higher profit. It is important to note that, in expectation, Single Sourcing from China in this game would result in higher profits than the TBS policy. However, our analytic model of the TBS does not capture dynamic control from China, since it essentially assumes the lead times are too long to allow for closed-loop control. Therefore, a single sourcing policy that does allow for dynamic control may do better than the TBS policy. (Also note that the significant differences between the expected value of a policy and the profit under the specific realization of demand are due to the high volatility of demand.)

5. *Summary and Extensions*

This section describes how we conclude the class and summarizes the contributions of the game and possible extensions, as well as student reactions.

5.1. Summarizing key take-away's of the game

We end the class by summarizing the key insights from the game as:

1. Dual sourcing is a surprisingly complex management problem with potentially high rewards. The first step in determining an effective dual sourcing policy is to consider its entire source-to-destination cost. While it is straightforward (yet tedious) to compute most components of this total landed cost, the impact on inventory and thus working capital is important yet difficult to assess under dual sourcing.
2. Effective dual sourcing uses suppliers with very different strengths. In this game, we combined a low cost (but slow) supplier with a fast (but expensive) source. Effective dual sourcing policies are also tailored to the strengths of each source. The tailored base-surge policy has a standing order with the low cost supplier (allowing that supplier to level the workload and reduce costs further) and orders from the fast source only to react to demand surges. Such a policy allows the firm to reduce total cost by placing the majority of orders to the low cost supplier while guaranteeing high service levels by placing occasional, small orders from the responsive, yet expensive supplier.
3. While determining optimal dual sourcing policies is computationally complex (and requires simulation), it is valuable—and often sufficient in practice--to have some guiding formulae to support dual sourcing decisions. The square-root formula provides a simple starting point when determining strategic sourcing allocations.
4. Successful dual sourcing often requires having a strategy and sticking to it. Having a suboptimal policy is typically better than continuously changing the decision process. For example, a strategy should allow the manager to react to high demands by increasing orders; however, reducing the target inventory level (base-stock) when observing a few low demand periods usually leads to low performance.

The game also provides a background to illustrate the benefits that companies have experienced after transitioning from sole to dual sourcing, including:

- Significant air freight reduction – by moving to dual sourcing some companies have seen 70%+ reductions in air freight.

- Faster response to changes in demand, as opposed to missing demand. Under single sourcing from a low cost source, by the time a demand change was recognized, it was too late in the season to respond.

- Improved ability to manage new product introductions. When products were sole sourced from an offshore low cost supplier, 6 to 9 months of product were in the pipeline before the company recognized poor sell-through which left them with significant obsolete inventory.

5.2. Possible Customizations and Extensions of the Game

The game can easily be extended along several dimensions, in order of suggested importance and ease of implementation:

- Customization: the game software currently allows the instructor to change the names and locations of the sources. For example, the fast source can be domestic and two locations could be France and Poland. The instructor can also modify the demand realization and any other parameters.
- Operational extensions: Changing the backlogging assumption in the game and allowing for lost sales is a straightforward modification to the game. So is adding a physical holding cost to the current financial opportunity cost or a fixed ordering cost. In addition, it is easy to modify the software and allow for stochastic lead times (but we found that this makes the game much harder to understand for students without offering any new insights).
- Financial extensions: It is straightforward to have a higher interest rate for borrowing than saving. One can also introduce a default threat by adding a constraint on maximal borrowing.
- Non-stationary parameters: For example, instead of assuming i.i.d. demand, the demand distribution could reflect the different stages of the product life cycle of introduction, growth, maturity, and decline. The same can be done with prices, costs, exchange rates, and interest rates.
- Competitive market: For example, any demand unmet by one team spills over to the other teams. This adds another layer of complexity, even for a simple single-period, single-source model where one can choose different excess demand splitting rules; see Lippman and McCardle (1997).
- Design-marketing-sales decisions: For example, teams can boost demand by spending money on marketing campaigns and sales incentives. This would necessarily imply that teams no longer observe a common demand. Another extension could be to include product and supply network design, similar to the *Global Supply* Chain *Management* Simulation of *Enspire Learning*.[5] While that simulation focuses on the design of the products and the supply network, it allows the players to modify the sourcing allocation in a limited manner. In contrast, our game is simpler and focuses on managing an existing supply network in a dynamic setting,
- Multiple items or multiple stages: while we do not suggest this, one could extend to multiple products or to a multi-stage supply chain along the lines of the famous Beer Game. In essence, this would lead to a Beer Game with dual sourcing.

[5] http://www.enspire.com/simulations/gscms. Also distributed by Harvard Business Publishing under Prod. #: 6107-HTM-ENG.

The game can also be used to conduct experiments to test research hypotheses. For example, teams may stick closer to their original plan and policy if they experience the game asynchronously or are not informed of other teams' bank balances.

While this shows how the game can be extended, we strongly advise to play the game in its current simple format because it allows students to map their decisions to subsequent results and thus enhances the learning from the game.

5.3. Instructor Experience so far

So far, the authors have played the game in about 10 business school class settings: with full-time MBA students in an operations strategy elective; with faculty members at a research conference; and in several executive education courses. Next we will play it with our doctoral students. We believe the game would also work well with graduate, as well as advanced undergraduate, students in engineering.
The response has been uniformly positive. Executives appreciate the realism of the game and the complexity inherent in dual sourcing. After playing the game using only intuition, the class enhances their perceived value of the academic research and simple guideline formula. Students often suggest that the game should be played with a cross functional team, representing their marketing, sales, financial, and operational groups. Such experience would convey the importance of inter-functional coordination and collaborative forecasting.

From an instructor's perspective, the game is easy to explain, has minimal requirements (it needs only a few laptops connected to the Internet), and is simple to set up. A single instructor can easily run the game with as few as 5 students or as many as 60 students without needing additional assistants or diminishing the effectiveness of the game. The game not only successfully achieves the pedagogical objectives; it also highlights the value of academic research for a realistic and important business problem.

Acknowledgments

We are grateful to Cort Jacobi and Ruchir Nanda of Deloitte Consulting for ongoing collaboration and insightful discussions on dual sourcing.

References

Allon, G. and J. A. Van Mieghem (2009). "Global Dual Sourcing: Tailored Base Surge Allocation to Near and Offshore Production." To appear in *Management Science*.

Lippman, S. A. and K. F. McCardle (1997) The Competitive Newsboy. *Operations Research* 45(1)54-65.

Van Mieghem, J. A. (2008). *Operations Strategy: Principles and Practice*. Dynamic Ideas, Belmont, MA.

Veeraraghavan, S. and Scheller-Wolf, A. A. (2006) Now or Later: A Simple Policy for Effective Dual Sourcing in Capacitated Systems *Operations Research* 56,; 850-864.

Janssen, F. and de Kok, T. (1999) "A two-supplier inventory model*", International Journal of Production Economics*, 59(1-3); 395--403.