

How to Win “Spend” and Influence Partners: Lessons in Behavioral Operations from the Outsourcing Game

Jason Amaral

Emeraldwise, LLC, Woodside, California 94062, USA
jamaral@emeraldwise.com

Andy A. Tsay

OMIS Department, Leavey School of Business, Santa Clara University, Santa Clara, California 95053, USA
atsay@stanfordalumni.org

The Outsourcing Game is a role-play simulation that has been deployed in industry and academic training courses worldwide. It incorporates the concepts of hidden actions, hidden information, and misaligned incentives, and conveys messages about power, trust, and reputation. The game depicts the adventures of Acme, the brand owner of a product manufactured by an outsourced supply chain. Through a series of negotiations, Acme attempts to influence its partners (two suppliers and two service providers) by distributing its procurement “spend.” These partners, in turn, sway each other via side payments. To simulate the non-linear shifts in power that occur as outsourcing increases, we represent decision-making by a voting scheme with uneven vote allocations. This paper analyzes a database of game results to reveal behavioral factors that can undermine conspicuous win–win process improvements. For instance, preferences can be sensitive to the sequence in which the alternatives are encountered; decision-makers might value not only their own rewards, but also fairness in the allocation of total gains; and effectiveness of negotiation tactics will vary with community norms of acceptable behavior. Along the way we extend the political economics literature about power in block-based voting by proposing a heuristic approach for incorporating voter preferences.

Key words: supply chain management; procurement; negotiation; simulation; cognitive biases

History: Received: August 2006; Accepted: November 2008 by Panos Kouvelis; after 1 revision.

1. Introduction

Many businesses have embraced outsourcing as a way to focus on “core competencies,” increase flexibility, and improve return on assets, among a long list of aspirations. Yet many have also misunderstood or underestimated the costs and risks that arise from converting internal functions into services procured from outside firms (cf. Tsay 2009a). The body of literature documenting these problems has grown rapidly in recent years (e.g., Amaral et al. 2006, Barthelmy 2003, Doig et al. 2001, Earl 1996, Lakenan et al. 2001, Thurm 2007). Our concern in this paper is with how outsourcing complicates decision-making once power is distributed across a constellation of autonomous firms whose relationships are shorter term and more transactional.

While these issues are salient for any business venture, we focus on the endeavor of stewarding a product from concept to market, and then operating the resulting supply chain. Inspired by success stories in manufacturing outsourcing, companies have rushed to outsource other operations as well—prod-

uct design, parts procurement, logistics, inventory management, product take-back and repair, supply chain IT, etc. (cf. Bardhan et al. 2007, Tsay 2009a,b). The management of the resulting extended enterprise thus takes on an increasingly political flavor, giving new importance to diplomacy in situations where central authority was once adequate. Here supply chain managers are finding that their destiny depends heavily on negotiation prowess, and uneasily grapple with an environment that includes the prospect of deception, or at least omission, by their business “partners.”

To convey such lessons to current and future managers, we created a teaching simulation called “The Outsourcing Game” (OsG) (Amaral and Tsay 2008). Because many of the crucial elements of the real setting can be emotional and contextual, we felt that an experiential exercise would be the most effective vehicle. To elicit participant behaviors consistent with the hazards uncovered in our academic and industry research (cf. Amaral et al. 2006), we designed a realistic, but simplified, outsourcing scenario that

combines hidden information, misaligned incentives, and hidden actions with distributed power and weak supply chain governance.

The OsG was introduced in 2003 in an MBA elective at Santa Clara University, and has been run in more than 35 sessions with nearly 1000 participants (in MBA and executive courses at Arizona State University, IMD, MIT, Nyenrode Business Universiteit, The Ohio State University, Purdue University, Santa Clara University, Tulane University, University of Utah, Washington University in St. Louis, and The Wharton School, and as part of a day-long workshop delivered with iSuppli Corporation and at domestic and international sites within Hewlett-Packard).

The OsG has been received well by both industrial and academic audiences. According to a procurement veteran with more than 20 years experience at large, multi-national corporations,

The Outsourcing Game is based on the reality of dealing with suppliers and their motivations. I've read and studied the process of contracting, and have practiced it for many years at large corporations. However, compared with a stereotypical training session, the game demonstrates that outsourcing is a lot more than just costs, margins, and negotiations. It helps you experience personally the drivers and motivations that will make or break an outsourced relationship.

One professor remarked,

The students thoroughly enjoyed this fast-moving simulation that embraces the complexities of managing an outsourced supply chain—including a consideration of the needs and concerns of the various parties involved. The broad perspective provides critical insights into supply chain power and relationship issues.

The OsG was a key part of a teaching portfolio that was honored with the Production & Operations Management Society's Wickham Skinner Award for Teaching Innovation for 2007.

While our objectives for developing the OsG were pedagogical, we have documented the results from 22 sessions over 4 years with the intention of conveying to researchers and managers some intriguing observations. This is with the caveat that we did not conceive the OsG to be a research platform and did not collect our data under perfectly controlled conditions. Indeed, the realism required to establish credibility with industry-savvy participants and maximize their learning necessitates that the game have more “moving parts” than is usually compatible with academic research protocols in psychology, experimental economics, and negotiation. Specifically, our framework includes five independent parties (three is

more common in behavioral studies of coalition-building) with highly imbalanced power, payoffs that can be transferred among players in almost unlimited ways, and information that is almost entirely privatized to the individual parties (behavioral studies typically keep hidden only very few items of information). This realism is both a weakness and a strength of our data set.

Thus, we will not attempt to prove or disprove any hypotheses about behavior in outsourced environments, or make any sharp claims about statistical significance. Our contribution will be to document where our experiences with the OsG conflict with or support the assumptions and predictions of the extant body of theory and experimental findings (cf. Bazerman et al. 2000, Bendoly et al. 2006, Gilovich et al. 2002, Kagel and Roth 1995). This can motivate more targeted investigation of the observed behaviors that participants have confirmed to resonate with their experiences in industry. Along the way we will connect to and extend the political economics literature about power in block-based voting by proposing a heuristic approach for incorporating voter preferences.

The remainder of this paper is structured as follows. Section 2 positions the OsG as an experimental protocol. Section 3 describes the resulting data set. Section 4 presents the major outcomes, and why they run counter to intuition. Sections 5 and 6 link the financial outcomes to the firms' power and negotiating postures (as perceived by the other parties), respectively. Section 7 concludes. A brief Appendix resolves a technical obstacle to the application of the Banzhaf Power Index (BPI) to the OsG.

2. Summary of the Game

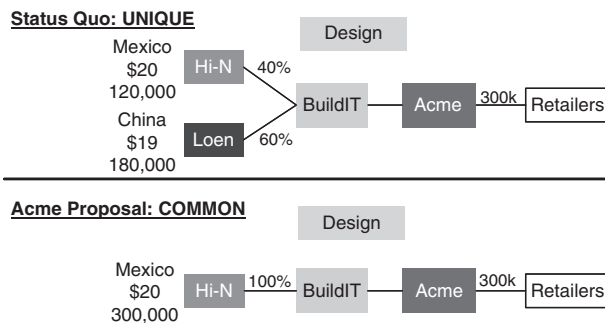
Due to space limitations, this paper will provide only the OsG details necessary to frame the game as an experimental protocol and communicate the behavioral findings. A full presentation of the execution of the OsG as a teaching instrument appears in Amaral and Tsay (2008).

Game participants are divided into five teams. One represents *Acme*, a company selling a branded product with high-end and low-end versions that are differentiated by a subassembly. The other four teams are *Acme's* supply chain partners: *BuildIT* (contract manufacturer), *Design* (outsourced design firm), *Loen* (supplier of low-end subassembly), and *Hi-N* (supplier of high-end subassembly).

The game centers on *Acme's* desire to switch to a common subassembly, as illustrated in Figure 1. In this paper we will refer to the status quo and *Acme's* proposal as *UNIQUE* and *COMMON*, respectively.

Loen and *Hi-N* each earn a per-unit margin for supplying their respective components. *BuildIT* earns

Figure 1 The Decision Problem



a 5% manufacturing fee based on the average material cost, and also incurs expenses connected to the complexity of the product line (which would be reduced by a switch to the common subassembly). Design earns a \$20,000 bonus from Acme per dollar of reduction in the average material cost (and is comparably penalized for increases). Acme faces the other side of each of these financial flows, and also bears inventory expenses related to materials costs and lead-times. Design must pay for any product redesign. These financials are highly private information in that each firm knows only its own inbound or outbound financial flows.

Table 1 details how moving from UNIQUE to COMMON would change the financial flows. The

numbers are imposed on the teams by exogenous specification of payoffs under the two options. This simplifies the analysis for participants, as they are shielded from any complex functional relationships that might be needed to quantify these effects in practice. The financial gains or losses presented in Table 1 indicate that two of Acme's partners initially favor COMMON, but two oppose it.

To simulate the non-linear shifts in power that occur in practice as the scope and scale of outsourcing increase, we represent decision-making by a voting scheme with uneven vote allocations. One hundred total votes are allocated in the following fashion: Acme (42), Design (40), Loen (8), Hi-N (8), and BuildIT (2). A simple majority of 51 or more votes is required to implement COMMON. The vote allocation is our approximation of the power balance in a typical real supply chain, with certain parties in dominant but not omnipotent roles. With two significant exceptions (any Acme-Design partnership and an Acme-Loen partnership for UNIQUE), control of the supply chain requires building a coalition of at least three parties. As will be discussed later, this gives certain parties power that is disproportional to their endowment of votes.

The financial incentives and the allocation of votes make COMMON appear to be a *fait accompli*. It could be achieved through a "Grand Coalition" since the immediate beneficiaries of COMMON (Acme, BuildIT, and Hi-N) could fund unanimous support for that

Table 1 Financial Impacts of Changing from UNIQUE to COMMON

Team	Role in supply chain	Advantages of COMMON	Disadvantages of COMMON	Net financial impact from COMMON
Acme (42 votes)	Product brand owner serving markets in the United States	+\$450,000: Reduced safety stock, cycle stock, and pipeline inventory +\$12,000: Penalty paid by Design	– \$180,000: Increased material costs paid to Hi-N – \$9,000: Increased manufacturing fee paid to BuildIT	+\$273,000
Design (40 votes)	Outsourced design firm based in the United States		– \$50,000: Engineering expense to redesign low-end product – \$12,000: Penalty paid to Acme	– \$62,000
Loen (8 votes)	Supplier of low-end subassembly with production facility in China		– \$360,000: Lost margin on Acme business	– \$360,000
Hi-N (8 votes)	Supplier of high-end subassembly with production facility in Mexico	+\$180,000: Additional margin on Acme business		+\$180,000
BuildIT (2 votes)	Contract manufacturer with production facility in Mexico	+\$9,000: Increased revenue from Acme +\$21,000: Reduced manufacturing complexity		+\$30,000
Net Financial Benefit of COMMON, including Loen				+\$61,000
Net Financial Benefit of COMMON, excluding Loen				+\$421,000
Net Financial Benefit of COMMON, excluding both Design and Loen				+\$483,000

outcome:

\$450k in inventory savings + \$21k in complexity
reduction savings – \$50k compensation for redesign
cost – \$360k to Loen for lost margins
= \$61k improvement over UNIQUE.

Sharing this surplus equally, if desired, would make each of the five teams \$12.2k better off than in the status quo of UNIQUE.

An obvious “Majority Coalition” also favors COMMON, with Acme (42), BuildIT (2), and Hi-N (8) collectively owning 52 votes:

\$450k in inventory savings + \$21k in complexity
reduction savings + \$12k Design performance penalty
= \$483k improvement over UNIQUE.

Sharing this surplus equally among Acme, BuildIT, and Hi-N would make each \$161k better off than in the status quo.

If inclusiveness were valued, some of the surplus could be used to coax Design into another majority coalition for COMMON. Loen would be cut out of that supply chain, leaving a total benefit of \$421k for the others to share (\$105.25k each, if shared equally).

In the world of theory, once a win-win outcome is identified, the assumption is that it will naturally prevail. The reality, as faced by OsG participants, is that the partnerships must be painstakingly built before anybody can pocket a single penny of benefit.

The interactions among teams are structured as follows. In Round #1 each team evaluates its private information and then chooses one of the options in Vote #1. This compels each team to understand the problem structure and evaluate the two options, establishing the initial landscape of preferences. Consistent with the financial terms described above, COMMON typically prevails here 52–48, with Acme, BuildIT, and Hi-N voting in favor (82% of sessions in our data set). This shifts Design and Loen into crisis mode.

In Round #2, the teams meet in two round-robin sets (“preliminary” and “final”) of one-on-one talks, with each meeting lasting five minutes. The negotiations conclude with 15 minutes of multi-party meetings as needed. One objective for each team during these meetings is to sign contracts that commit other teams to support a favored option (or conversely to extract compensation in exchange for supporting other teams). Another objective is to obtain favorable relationship scores from the other teams (described below). Vote #2 follows these talks.

This structure puts the OsG somewhere in between the categories of cooperative and non-cooperative games. Unlike, say, the classical Prisoner’s Dilemma

from non-cooperative game theory, during these talks the teams are allowed to communicate with other teams, and reveal information and share payoffs as they see fit. However, the OsG is more complex than what has been successfully analyzed in cooperative game theory, due to the freedom allowed in the sharing of information and funds. All forms of bluffing and misdirection are allowed en route to signing contracts. One basis for bluffing is to invoke factors that do not actually appear in the payoff numbers. For instance, Loen can try to persuade Acme that multi-sourcing provides risk reduction or enables sharing of technical capabilities between suppliers. Loen can even claim that it will develop the capability to produce the HE subassembly, which is a complete fabrication given the assumptions of this game. Obviously, such behaviors would be nonsensical in games without information asymmetry.

Along with choosing to support COMMON or UNIQUE, each team also assigns a “strength-of-relationship” score to every other team, defined as follows:

- 1 = Strong negative: Mistrust; desire to “disengage” as soon as possible; seek to exploit
- 3 = Neutral
- 5 = Strong positive: High trust; desire to foster a long-term, win-win partnership

This metric is intended to link current actions to long-run goals, as the one-shot nature of the game might otherwise engender unrealistically myopic behavior. Participants learn that financial value and relationship scores will determine the winners, but are provided no relative weighting. We have no concrete sense as to how teams make the tradeoff between their explicit financial goals and their desire for certain relationship scores, although some teams have contracted with and paid rivals solely for relationship scores of “5.” However, the presence of this non-financial criterion further distances the OsG from the settings analyzed in existing theory.

If time permits, the OsG can be extended with Round #3 (another series of negotiations and a Vote #3). This round does not represent a second period of business. Instead it explores how the scenario might play out under a different set of conditions, which varied across the sessions in our data set. In some cases, the Acme team was given complete information, which ought to have ensured leverage for brokering the deals to achieve COMMON. In other cases, the default financial flows were presented to all teams, which exposed any prior “mistruths” and simulated a scenario of open-book accounting. Before the new negotiations, the teams are instructed to forget all preceding events (although the game relies on their

inability to do so, bringing in issues of trust and reputation). By this point, the teams are familiar with the rules and each other, so preliminary talks are no longer required. Each team meets with every other only once for 5 minutes, followed by 15 minutes of multi-party meetings. Like Round #2, Round #3 ends with teams assigning relationship scores, turning in signed contracts, and casting Vote #3.

The OsG concludes with a debriefing by the facilitator. This includes an analysis of financial results and relationship scores.

3. Game Outcomes

Does the existence of a Pareto improvement, especially one with no overt barriers to implementation, mean that it will naturally come to fruition? The OsG suggests that this is not guaranteed under realistic conditions with real human agents.

The results in this paper are based on 22 sessions with 497 participants conducted by the same lead facilitator (coauthor Amaral) between March 2004 and June 2007 (see Table 2). The majority of sessions were conducted in MBA/MS programs (64% of sessions and 74% of participants), and the rest were in industry settings (working professionals from supply chain, design, procurement, and sales functions). Each session included five teams in “cooperation.” Each student team had an average of 5.3 members and each industry team had an average of 3.2 members.

The primary outcome of the OsG is whether the supply chain can implement the Pareto-improving COMMON plan, as opposed to being left with UNIQUE. These are the results of Votes #1, #2, and #3.

Vote #1 is merely a straw man that establishes preferences and gives participants context for the subsequent negotiations. The meaningful votes are the ones taken after negotiations, i.e., Votes #2 and #3. Our data set covers 39 of these votes coming from 22 game sessions, of which five did not have time for Vote #3. Ignoring the differences between these two rounds of voting, and the rule modifications made as the OsG evolved over time, Table 3 indicates that COMMON was the outcome in only 1/3 of the votes.

Although the OsG structure and process has been fairly stable from the beginning, some modifications

Table 3 Frequency of Votes that Achieved COMMON

Sessions	Frequency of COMMON (out of Votes #2 and #3)	Variations in the team objectives
1 and 2	0/2 = 0%	• Net advantage of COMMON = \$1 K
3 through 10	5/16 = 31%	• Personal motivation for Loen to minimize losses • Net advantage of COMMON = \$1 K
11 through 22	8/21 = 38%	• Personal motivation for Loen to minimize losses • Net advantage of COMMON = \$61 K
Overall	13/39 = 33%	

were made along the way to improve the learning experience. Three are discussed below.

- *Loen gained an additional decision objective.* Since Loen stands to lose its entire profit of \$360k if COMMON is adopted, paying out any less would be viewed as an improvement. Indeed, in sessions 1 and 2 Loen surrendered almost everything (\$350k and \$290k, respectively) to preserve UNIQUE (neither session had a Vote #3). As such extreme behavior detracted from the OsG’s message, as well as being somewhat unrealistic, we altered Loen’s incentives to encourage preservation of financial assets. Starting in session 3, Loen’s private information included the following personal motivator: “*The outcome of this negotiation could make or break your career. Based on heated internal discussions, you believe the personal impact to you would be . . . ‘Get fired’ (for paying out > \$360k), ‘Get demoted’ (for paying out between \$360k and \$320k), ‘Keep job’ (for paying out between \$320k and \$240k), or ‘Get promoted’ (for paying out < \$240k).*” From that point forward, in all but two rounds Loen only signed contracts that would lead to a promotion. Because the other teams typically remain unaware of this incentive structure, highlighting it during the debrief discussion underscores the importance of differentiating between the firm-level objective and the *personal motivations* of the relevant individuals.
- Relationship scoring was introduced. This was added in session 7.
- *COMMON was made even more financially superior to UNIQUE for the system.* In the beginning, COMMON was only superior by \$1k. Starting in session 11, COMMON’s advantage became \$61k due to a reduction in Design’s redesign expense (from \$100k to \$50k) and an increase in BuildIT’s benefit from complexity reduction (from \$11k to \$21k).

Although the shift in COMMON’s financial advantage starting in session 11 increased the frequency of

Table 2 Participants and Sessions

Environment	Sessions of OsG	Participants	Average participants per session (# per team)
MBA/MS program	14 (64%)	370 (74%)	26.4 (5.3)
Industry workshop	8 (36%)	127 (26%)	15.9 (3.2)
Total	22 (100%)	497 (100%)	22.6 (4.5)

COMMON outcomes (from 31% to 38% of votes), it did not significantly reduce the relative value obtained by Loen. Figure 2 shows the allocation of value received by each team depending on vote outcome and system-wide financial attractiveness of COMMON. Throughout all the game modifications, COMMON remained elusive. The next section evaluates why.

4. Why Did “Win–Win” Not Win?

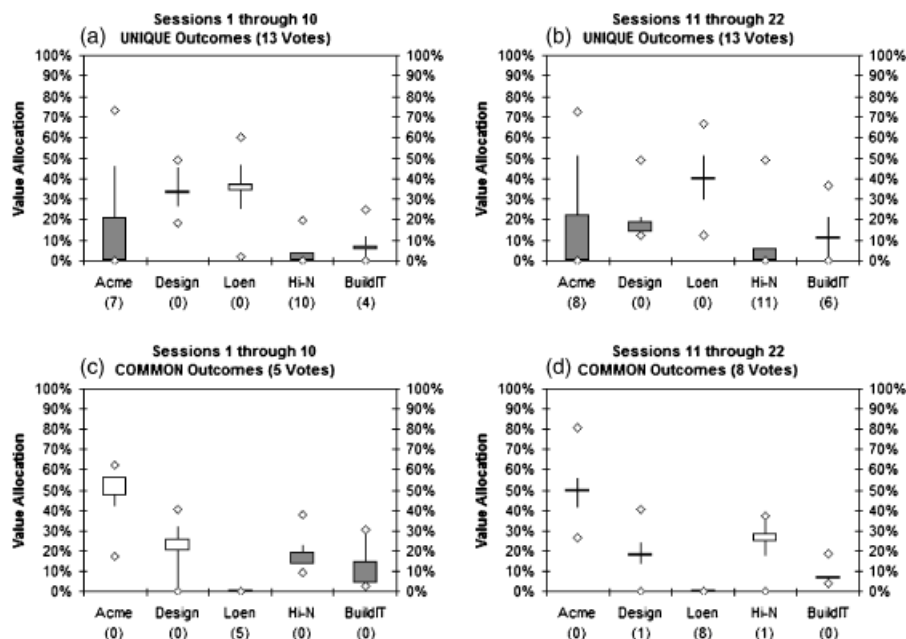
Rapoport (1970) summarized two of the central questions in coalition research: Which coalitions are likely to form? How are the gains from cooperation likely to be shared? This vast body of work includes game-theoretic analyses of supply chain settings (see Nagarajan and Sosis 2008 for a recent review), some attempting to predict the outcomes of sequential negotiations among more than two parties (e.g., Marx and Shaffer 2001, Nagarajan and Bassok 2008). However, such models are typically highly stylized, and presume perfectly rational behavior. They assume that (almost) all parameters are common knowledge, which rules out the kinds of bluffing and misdirection that are central to the OsG (and real life). One benefit of juxtaposing these models with our setting is to underscore the extent of simplification of assumptions that is needed to obtain analytical results. Because the extant prescriptive theory seems silent

on the case of multiple rounds of sequential negotiations with private information, not to mention a distribution of power based on ownership of votes, here we simply present our data, and do our best to describe the behaviors that emerged. We will postulate plausible root causes. Along the way we will identify behavioral phenomena that are consistent with existing literature, and articulate interesting conjectures that could be tested in focused experiments.

4.1. Grand Coalition for COMMON

As noted in Section 2, the net financial advantage of COMMON over UNIQUE could be used to fund a grand coalition. The quickest way to achieve this would be for all teams to reveal their private financial information, and agree to an equal division of the proceeds of collaboration. The rules do not prohibit this kind of collaboration. Instead, each team’s private instructions contain the phrase “*As part of negotiations, you may decide to share some or all of your private information with other organizations.*” Furthermore, the game participants typically already know each other socially, either as fellow students or co-workers. Laboratory experiments have found that people act most “fairly” in face-to-face encounters with people from common social circles (cf. Camerer and Thaler 1995, Cialdini 2007, Roth 1995). Yet in spite of these promising conditions, we have never witnessed a five-party coalition for COMMON.

Figure 2 Value Allocations by Team and Vote Outcome



Diamonds in each vertical column indicate the maximum and minimum value allocation observed. Tops and bottoms of “whiskers” indicate third and first quartiles. Tops and bottoms of gray boxes indicate the mean and median values, respectively. This is reversed for white boxes, where tops and bottoms indicate median and mean values, respectively. Values in parentheses below team names indicate the number of votes in which the team ended with a 0% value allocation. In panel (c), Design twice received a nonzero but trivial value allocation (0.25%) based on a minor reduction in the HE material price (and a corresponding reduction in its performance penalty).

In fact, the OsG contains features that are known to psychologically discourage this open-book cooperation. Experiments have found that private information, asymmetric information disclosure, and a lack of knowledge about “who knows what” result in more frequent disagreements and uneven payoff distributions (Roth 1995). Even the nomenclature matters, as in the finding of Liberman et al. (2004) that simply calling an N-Move Prisoner’s Dilemma “The Wall Street Game” instead of “The Community Game” reduced by more than half the first round cooperation (from 71% to 33%) and overall cooperation (from 66% to 31.5%). In the OsG, “Us-versus-Them” mindsets result from dividing and physically segregating the participants into teams with memorable names. Each team’s private instructions also contain a confidentiality notice and the sentences “*The details below on costs and benefits are private and are not known to any other group. . . . Bluffing is allowed.*” Of course, these factors likely exert a disruptive influence on all nascent coalitions, not only the grand coalition.

4.2. {Acme, Design} Coalition for COMMON

In spite of the psychological obstacles discussed thus far, Acme could realize that a bilateral partnership with Design would immediately lock up 82 votes. Because Design’s initial opposition to COMMON is due to factors entirely under Acme’s control—in the performance penalty for materials cost increases and in Design’s obligation to absorb redesign costs—Acme could end the game instantly by voiding these terms. This tactic certainly does not require much imagination. Were Design an internal R&D organization, netting out all of the direct costs and benefits would be natural. Also, the dysfunctional nature of Design’s financial incentives is immediately obvious to most people. Yet Acme and Design signed a contract with each other in only 23% of Rounds #2 (five out of 22), which we attribute to three main drivers.

First, Acme might feel blameless for Design’s situation. Acme members did not *personally* create the current incentives and costs, so might be inclined to operate within the legacy framework much like a newly promoted manager would. Further, Acme’s instructions contain the sentence “*After talking with your lawyers, you believe that despite their complaints, Design would be contractually bound to do the redesign as part of the retainer fee.*” Both of these explanations are consistent with psychological studies showing a natural deference to authority without conscious deliberation (cf. Cialdini 2007).

Second, any COMMON outcome in Vote #1 psychologically positions the two teams as antagonists going into subsequent rounds. In 18 out of 22 sessions, the first vote went purely along economic lines with Acme, BuildIT, and Hi-N forming the victorious

COMMON coalition. Notably, in 19 out of 22 of the Votes #1, Acme chose COMMON and Design chose UNIQUE. Design was frequently in the *losing* coalition (17 out of 22), while Acme was frequently in the *winning* coalition (20 out of 22 sessions). So Vote #1 could intensify Design’s feelings of impotence and defensiveness, while giving Acme a greater sense of control and possible overconfidence (making Acme’s preferred outcome seem more likely per the availability heuristic; Tversky and Kahneman 1974). Acme might also develop a sense of entitlement towards its payoff from COMMON (\$273k) (per the anchoring heuristic; Tversky and Kahneman 1974). Any concession to Design would then be perceived as a loss from its “endowment” even though it might still be an improvement over the status quo (per Prospect Theory; Kahneman 2003, Kahneman et al. 1991).

Third, and perhaps most counterintuitively, the negotiations *start* with a meeting between Acme and Design. In coalition-building, the sequence of negotiations is known to influence their outcome (Sebenius 1996, 2004). One reason is that negotiators are typically reluctant to make early commitments before exploring their other options (Lax and Sebenius 2006, Susskind 2004). So Design will be susceptible to subsequent proposals from other parties (particularly Loen). Marx and Shaffer (2001) and Nagarajan and Bassok (2008) both showed mathematically that the sequence of negotiations affects the outcomes in supply chain bargaining games, though not in the psychological way that we postulate here.

The relationship scores suggest that the {Acme, Design} relationship took on an adversarial tone for reasons consistent with the preceding discussion. Compared with the average scores given to the other three teams, Acme rated Design 1.0 lower, while Design rated Acme 0.2 points lower. Also, Acme did not internalize that *it* was the cause of relationship tension, as Acme scored Design 0.8 points lower than Design scored Acme (2.3 versus 3.1). To put these in context, compare with the average scores given by the other three teams. Design actually thought more highly of Acme than the others (3.1 versus 3.0), but Acme rated Design lower than the others (2.3 versus 3.0). The relationship scores will be explored more comprehensively in Section 6.

4.3. {Acme, Hi-N, BuildIT} Coalition for COMMON

Acme, BuildIT, and Hi-N together have sufficient votes (52) to achieve COMMON. Absent other agreements, such a vote would provide a combined benefit of \$483k, including the \$12k performance penalty paid by Design. An equal division would give \$161k to each (or \$157k each under the conditions of sessions 1 through 10), dominating the status quo

baseline of \$0 and the grand coalition surplus of \$12.2k. (In addition, an inventory savings of $(1 - \text{HE}/\$20) * \400k will result if Hi-N discounts the HE sub-assembly. If Hi-N reduces the price to match its status quo margin (\$19.40/unit), the COMMON coalition can still share \$483k because the inventory savings exactly offset the elimination of Design's performance penalty. The "vertically integrated" COMMON coalition could share \$503k if the HE is sold at cost (\$19/unit).)

However, any plan that assumes an equal division of gains can fall apart if any party feels entitled to more (Raiffa 1982). We found that discussions among these three teams typically got bogged down over the correct "fair share benchmark" (Roth 1995). Certainly each team can be thought to bring a different share of financial value to the table: 57% from Acme, 37% from Hi-N, and 6% from BuildIT. Per the anchoring heuristic mentioned previously, Acme and Hi-N are likely to feel ownership for their respective contributions.¹ Alternatively, the teams might fixate on the number of votes they contribute to the winning block: 81% from Acme, 15% from Hi-N, and 4% from BuildIT. Finally, BuildIT might argue for an equal share since its 2 votes are equally vital to the winning coalition.

Any three-way division of the \$483k surplus from COMMON would leave at least one of three members with no more than \$161k of benefit (this is an example of the Pigeon-Hole Principle). Such a member represents a vulnerability in the coalition, since Loen's manager can pay out up to \$240k to "divide and conquer" and still earn the promotion. BuildIT is the natural target since its mere two votes give it the weakest claim on the surplus generated by COMMON.

Such a dynamic seems to have occurred. When BuildIT participated in a COMMON coalition, it received approximately 12% of the surplus. However, of the 39 Votes #2 and #3, 26 led to UNIQUE (67%). This was most frequently (11 times out of 26, or 42%) due to the coalition of {Design, Loen, BuildIT}. Overall, BuildIT was involved in 69% of the winning UNIQUE coalitions (18). Excluding the first two sessions of our data set, Loen and Design were able to achieve a UNIQUE outcome by paying less than the "equal COMMON split" (\$157,667 or \$161k) in 63% of the votes (15 out of 24).

5. Predicting the Financial Outcomes from Vote Endowments

At its core, the OsG is a set of negotiations through which teams buy and sell the blocks of votes with which they have been endowed. The key research question is how relative power (as embedded in the initial allocation of votes) translates into financial outcomes.

The literatures of law and political economics contain various methods of quantifying relative power in block voting systems, for example the Shapley and Shubik (1954) and Banzhaf (1965, 1968) power indices. (The Banzhaf concept was actually articulated earlier by Penrose 1946.) For a focused tutorial on the mathematics of these two, see Bowen (2001) and Livingston (2003). These two indices use similar logic, but the former makes distinctions for the sequence in which members join coalitions whereas the latter is concerned only with the coalition roster. Hence, we favor the BPI as the starting point for analysis of the OsG (although the two indices produce similar conclusions here). While there is controversy over whether the BPI's assumptions lead to the most meaningful measure of power, this index is a well-accepted construct whose technical attributes have been studied extensively (cf. Dubey and Shapley 1979) and that has been extended in many ways.

The BPI reports a power value for each voting block, in the context of the full allocation of votes and the number of votes required to win an election between two options. Coming to notoriety as part of the public criticism of the Electoral College system for selecting American presidents, the BPI focuses on the phenomenon of the "swing vote," i.e., the voting block whose defection would turn a winning coalition into a loser. A block's BPI is the frequency with which it appears among the set of "swing" incidents. This captures the notion that while controlling more votes always means more power (weakly), power is not directly proportional to the number of votes owned. For instance, if winning requires a simple majority and one player possesses 51 votes out of 100, all other players have zero power. But if, say, an election is conducted among three players with 49, 49, and 2 votes, respectively, the 2-vote player is as powerful as both others.

We are not aware of a closed-form formula for the BPI. The computation entails combinatoric enumeration when the number of blocks is small (which we will illustrate for the OsG below), and sometimes switches to Monte Carlo simulation or a Stirling's approximation when the number of blocks is large. Automated calculators are readily available online.

In the following calculations we abbreviate the OsG parties as A, D, L, H, and B to conserve space, and use shorthand, such as ADL to refer to the {Acme, Design, Loen} coalition.

STEP 1: Determine all coalitions that achieve 51 or more votes (for COMMON).

⇒ AD, ADL, ADH, ADB, DLH, ALH, ALB, AHB, ADLH, ADLB, ADHB, ALHB, DLHB, ADLHB

STEP 2: Determine the critical members (“swing voters”) in each of these coalitions (denoted by underlining).

⇒ AD, ADL, ADH, ADB, DLH, ALH, ALB, AHB, ADLH, ADLB, ADHB, ALHB, DLHB, ADLHB

STEP 3: Find the total number of instances of a coalition member being critical.

⇒ STEP 2 identified 26 swing-vote instances (the number of underlined letters).

STEP 4: Count the number of times each coalition member is critical.

⇒ A: 10; D: 6; L: 4; H: 4; B: 2

STEP 5: The BPI values are the numbers from STEP 4 divided by the number from STEP 3.

⇒ These are displayed in the first data column of Table 4.

A striking feature of these results is that BuildIT has much more power than its endowment of two votes might suggest. On the other hand, Design’s block of 40 votes is much less powerful than Acme’s block of 42.

One criticism of the BPI is its assumption that each block will vote by flip of a fair coin. Gelman et al. (2002) have proposed empirically incorporating voter preferences through historical voting data, essentially assigning a distinct weighted coin to each block. We require a different approach because the parties in the OsG enter without history, but with explicit (financial) incentives to favor one option over the other. To the best of our knowledge, the power indices in the existing literature are silent with respect to this setting.

To address this, we define a “preference-adjusted” BPI (paBPI). This is computed by first filtering out the implausible coalitions, and then applying the classic BPI logic to the remaining ones. This step resembles the process of solving for game theory equilibria by identifying and eliminating dominated strategies. Just as a given game might have no dominated strategies, there is no guarantee that any coalitions can be ruled out, in which case the paBPI and BPI values would match. The exclusions must be a rational consequence of the original game assumptions. This means that

while historical incidence rates might inspire the analyst to exclude a certain coalition, a logical argument is required. That is, we cannot conclude that a given coalition will never arise just because we have not seen it arise before.

In the OsG, the initial financial positions rule out the following coalitions:

1. The five-team coalition for UNIQUE. The teams would realize the dominance of COMMON and switch to this option.
2. A UNIQUE coalition that includes all parties except Design. Design has a natural incentive to join in.
3. Any coalition for UNIQUE that does not include Loen. Loen is the primary funding agent of the UNIQUE coalition.
4. Any coalition for COMMON that includes Loen, except the grand coalition. Loen financially favors UNIQUE far more than any other team.

Striking-through and italicizing the excluded coalitions produces the following list of remaining winning coalitions, with critical blocks underlined:

Winning COMMON coalitions: AD, ~~*ADL*~~, ADH, ADB, ~~*DLH*~~, ~~*ALH*~~, ~~*ALB*~~, AHB, ~~*ADLH*~~, ~~*ADLB*~~, ADHB, ~~*ALHB*~~, ~~*DLHB*~~, ADLHB

Winning UNIQUE coalitions: ~~*AD*~~, AL, ~~*AH*~~, ADL, ~~*ADH*~~, ADB, DLH, DLB, ~~*DHB*~~, ALH, ALB, ~~*AHB*~~, ADLH, ADLB, ~~*ADHB*~~, ~~*ALHB*~~, DLHB, ~~*ADLHB*~~

Here we need to articulate the list of all winning coalitions for both outcomes, since the exclusion cases affect each list differently.

The remainder of the paBPI follows STEPS 3–5 from the BPI algorithm. STEP 3 shows 23 distinct swing-vote scenarios (10 among the winning COMMON coalitions and 13 among the winning UNIQUE coalitions). STEP 4 shows the individual coalition members to be critical with the following frequency: A: 9; D: 6; L: 4; H: 2; B: 2. STEP 5 produces the paBPI values in the second data column of Table 4.

Even though Loen and Hi-N possess the same voting power (and hence have equal BPI values), the paBPI does not view them as equals. This is because the two parties have very different preferences.

Table 4 Power Indices and Predictions of Financial Outcomes

	Power index		Average final share of system financial value		
	BPI (%)	paBPI (%)	Sessions 1 through 22 (%)	Sessions 1 through 10 (%)	Sessions 11 through 22 (%)
Acme (42% of votes)	38.5	39.1	32	29	35
Design (40% of votes)	23.1	26.1	24	30	19
Loen (8% of votes)	15.4	17.4	24	24	23
Hi-N (8% of votes)	15.4	8.7	11	8	14
BuildIT (2% of votes)	7.7	8.7	10	9	10

Conversely, the paBPI views Hi-N and BuildIT as having comparable power in spite of a disparity in vote ownership. In the shadow of Acme, their natural but much more powerful partner, they both seem to have roughly the same clout.

The intent of this section was to predict how the financial value in the system would be distributed as a result of negotiations. Table 4 shows that for the 22 sessions in our data set the BPI and paBPI values both accomplish that much better than does the vote allocation.

The paBPI correctly estimates that Loen will do approximately twice as well as Hi-N on average, whereas the BPI (and, obviously, the initial vote endowment) views the two firms as equal. Both indices seem to overestimate Acme's financial outcome and underestimate Loen's, which might be attributable to the "psychological disadvantage" the game structure imposes on Acme *vis-à-vis* its initially strained relationship with Design.

Our paBPI should not be expected to have perfect predictive power since it incorporates only a subset of the initial financial endowment values. Also, any negotiations conducted by human players with bounded rationality are inherently unpredictable.

Table 5 reports the frequency of each winning coalition during real game play. Notably, some of the coalitions that were admissible for the paBPI never arose.

A close comparison of the winning coalitions in our data set and Figure 2 (summary of value allocation by team and vote outcome) might suggest several apparent inconsistencies. However, they can easily be explained. Acme was excluded from 16 UNIQUE co-

alitions, but received 0% value in only 15. In one vote, Loen paid Acme \$50k for a "5" relationship score. BuildIT was excluded from 8 UNIQUE coalitions, but received 0% value in 10. BuildIT voted UNIQUE twice although it did not enter into a contract or receive any payment. Design was excluded from three COMMON coalitions, but received 0% value in only 1. In two sessions, it faced a reduced performance penalty because Hi-N lowered the HE subassembly price.

6. Negotiation Behavior and Financial Performance

Since rule by fiat is largely a thing of the past in the modern outsourced supply chain, we have designed the OsG to shed light on how interfirm relations and economic rewards impact each other. Some insights emerge from correlating the relationship scores with the financial outcomes.

We acknowledge the intrinsic imprecision of using any single metric to assess the sum total of each party's tactics. Indeed, Thompson (2008) clarifies that the appropriate question when negotiating is not as black and white as "Should I be tough and competitive, or nice and cooperative?" Instead, there is a cooperative phase of "creating value" and a competitive phase of "claiming value" (Lax and Sebenius 1986) and each might merit a distinct negotiation approach. Further, if done appropriately (e.g., with consistent and transparent posture that favors "win-win" outcomes), "tough" behavior might not be perceived negatively from across the negotiating table. All these interpretations are culturally specific as well. With these caveats, here we interpret high

Table 5 Winning Coalitions in Our Data Set

Winning outcome	Winning coalition	Appearances during game play (out of 39 rounds)		
		Frequency	% of each outcome	% of all rounds
UNIQUE outcome in 26 rounds (67%)	ADLH (98 votes)	00:		
	ADLB (92 votes)	06: ■■■■■■	23.1	15.4
	ADL (90 votes)	04: ■■■■	15.4	10.3
	ALH (58 votes)	00:		
	DLHB (58 votes)	01: ■	3.8	2.6
	DLH (56 votes)	04: ■■■■	15.4	10.3
	ALB (52 votes)	00:		
	AL (50 votes)	00:		
	DLB (50 votes)	11: ■■■■■■■■■■	42.3	28.2
	AHB (52 votes)	03: ■■■	23.1	7.7
COMMON outcome in 13 rounds (33%)	AD (82 votes)	00:		
	ADB (84 votes)	00:		
	ADH (90 votes)	03: ■■■	23.1	7.7
	ADHB (92 votes)	07: ■■■■■■	53.8	17.9
	ADLHB (100 votes)	00:		

relationship scores to be indicative of niceness, fairness, and integrity, while low scores reflect unpleasantness, spitefulness, and untrustworthiness. The future incorporation of detailed surveys would help to confirm the motivations and impressions of participants.

To exclude the reputation effects that carry over into Round #3, here we consider only the relationship scores from Rounds #2. (Because most participants already knew each other previously, we could not completely eliminate all reputation effects.) For each team these data are grouped in three ways in Table 6: scores given by all other teams, scores given by those teams that voted the same way, and scores given by the teams that achieved contractual agreement (and, by extension, also voted the same way).

The mean relationship scores increase monotonically from left to right in each row (the direction of increasing economic ties), suggesting that the scores were assigned in a plausible and internally consistent way. Focusing on the teams one-by-one sheds light on how a team's endowment of votes and financial value translated into behavior, and consequently its standing in the eyes of the other teams.

As rated by voting peers and contract partners, Acme had relatively high relationship scores with a positive correlation with the value allocation it received. For Acme, niceness, fairness, and integrity seem to have paid off. Positive relations may have allowed Acme to capitalize on the strength of its voting power and financial endowment. It built and maintained winning COMMON coalitions at lower costs, or gained a favorable position in Loen's UNIQUE coalition.

Design's results were comparable to Acme's, although, across the board, the relationship scores were slightly lower and the correlations weaker. This is not surprising in that Design had nearly as many votes as Acme, but controlled less economic value. Positive relations allowed Design to leverage its options by either solidifying the alliance with Loen or overcoming the initial tensions with Acme.

Loen received the highest relationship scores of all. With everything at stake but only a middling degree of voting power, Loen used its financial resources to

buy votes. The high scores might be explained by the transparency of Loen's preferences. It was motivated to negotiate earnestly (based on usually losing in Round #1), and was understandably consistent in its desire for UNIQUE (based on the need to protect business). Moreover, the clarity of Loen's personal incentives provided an unambiguous reservation price, allowing it to make quick and tangible offers. In light of the time pressures of the game, as well as the hidden agendas and preliminary offers of the other teams, this attitude was likely to be viewed favorably. The lack of correlation occurred because all teams generally gave Loen high scores, whether or not they voted the same way, engaged Loen in a contract, or received a high payment.

Hi-N was somewhere in the middle of the ratings, which is consistent with its being somewhere in the middle in terms of voting power and control of financial value. As with Loen, Hi-N's preferences were more obvious—although Hi-N displayed a willingness to support UNIQUE for the right price. Hi-N's negative correlations were generally due to getting higher scores when it made higher payments.

BuildIT received noticeably lower relationship scores than did the other teams. Note that its rating by contract partners was not much better than the rating Loen received from all other teams. It also had strongly negative correlation values, indicating that BuildIT had to choose between "love or money." Notably, this tension was strongest in relationships with the ones targeting (or targeted by) BuildIT for contracts. Unlike Acme, BuildIT was dealing from a position of weakness, with only two votes and little financial value to offer. Its power came purely from being a swing-block of votes. Monetizing that power required playing hardball and inciting a bidding war. Demanding ransoms is surely not the way to make friends. But as Benjamin Franklin wrote in 1788, "it (is) more difficult for a man in want to act always honestly, as (to use here one of those proverbs) 'it is hard for an empty sack to stand upright'" (Lemisch 1961, p. 107).

In the OsG and the business world it attempts to simulate, the interdependence between interfirm

Table 6 Relationship Scores and their Correlations with Value Allocations for Vote #2 (16 Sessions in which Relationship Scores were Collected)

Team	As rated by all other teams		As rated by voting peers		As rated by contract partners	
	Mean (RS)	Corr (RS,FV)	Mean (RS)	Corr (RS,FV)	Mean (RS)	Corr (RS,FV)
<i>Acme</i>	3.0	0.3	4.0	0.2	4.2	0.6
<i>Design</i>	2.9	0.1	3.5	0.0	4.0	0.3
<i>Loen</i>	3.4	0.0	4.1	0.1	4.3	0.0
<i>Hi-N</i>	3.2	0.0	3.7	−0.6	4.3	−0.4
<i>BuildIT</i>	2.7	−0.4	3.1	−0.4	3.6	−0.6

RS, relationship score received; FV, value allocation accomplished via Round #2 activities.

relations and financial rewards are complex and situation-dependent. Consistent with the findings of negotiation researchers (Bazerman et al. 2000, Lax and Sebenius 2006, Thompson 2008), managers of outsourced supply chains cannot rely on simple aphorisms such as “nice guys finish first” (Dawkins 2006) or alternatively “it is much safer to be feared than loved” (Machiavelli 1515). Managers must deeply assess their situation, and investigate structural root causes when interpreting the seemingly antagonistic or less than forthcoming behavior of their supply chain partners.

7. Conclusion

Real-world supply chains are messy. They consist of numerous firms whose masses of employees have varying levels of experience and diverse cultural backgrounds. Hidden actions, hidden information, and misaligned incentives are ubiquitous, and outsourcing only exacerbates their impact. Real managers navigate this maelstrom with tactics better explained by bounded rationality and behavioral biases than by economic optimization.

At the other end of the spectrum are the supply chain models idealized to follow axiomatic order. These have been formulated and analyzed for decades, leading to very precise truths about very crude representations of reality.

Somewhere in the middle is a growing body of behavioral studies that also use stylized and highly structured simplifications of reality, but which include human subjects who are allowed to act human. Experiments carefully vary only one parameter at a time. Creativity in experimental design and adherence to strict delivery protocols are the coin of the realm.

Although first and foremost an educational instrument, the OsG is offered here in the spirit of this third body of research. Our experiments must be understood to have been closer to the messy real world than to the ideal, controlled laboratory. With that caveat, our findings offer a number of insights that can be useful to architects and managers of modern supply chains.

A message that emerges loudly from this work is that the mere existence of a conspicuous and significant win-win opportunity does not mean it will materialize. The OsG demonstrates the contrary, that successful realization of Pareto improvements might well be the exception rather than the rule. Any change requires a specific pathway, and moving the entire supply chain down this path is a matter of power and influence. Many would do well to heed the timeless advice from Dale Carnegie’s *How to Win Friends and Influence People*: “[T]he only way on earth to influence other people is to talk about what *they* want and show

them how to get it” (Carnegie, 1981, original copyright 1936, p. 33, italics in original).

The meaning of power in highly fragmented collectives is not intuitively obvious. Political scientists and economists have given this a great deal of thought, but we have encountered few managers who approach supply chain management as a political exercise in coalition-building. Far more common is to attack supply chain problems with an engineering mindset, focusing on those aspects that are more easily quantifiable and rationally understood.

In fact, our data suggest that supply chain managers, like all other managers, are susceptible to complications such as anchoring biases, sensitivity to the sequence in which potential partners are screened, or the need for the division of gains to seem fair. Ultimately we find that, time and time again, people will walk away from a nontrivial financial gain if this requires doing business with those they do not like or trust, even when any real risk of betrayal can be eliminated (in our case, by signed contracts) (cf. Fehr and Gächter 2000). This is not necessarily a bad thing, as the willingness of parties to incur personal expense to punish unsavory behavior may be a prerequisite for an orderly society and economy (Fehr and Gächter 2002). Nevertheless, a cynic might see this as “cutting off the nose to spite the face.” Our analysis helps resolve this conflict by establishing linkages between negotiation posture (“tough” vs. “nice”) and financial payoffs.

Such behavioral twists will resonate with the personal experiences of our readers, but are glaringly absent from the analytical model-based supply chain research. We hope this paper, based on a framework validated by numerous professionals with relevant real-world experience, can provide some motivation and structure to help remedy this shortcoming.

Acknowledgments

We thank the many instructors and participants all over the world who have enthusiastically embraced OsG. Their feedback has validated and enhanced the degree of realism in the role-play. We appreciate the efforts of the POM journal editorial team, which greatly improved this paper’s direction and exposition.

Note

¹For Acme, separate accounting for the \$450k inventory benefits may be justified. Unlike the other costs and benefits in the OsG, the inventory risk pooling benefits are *expected, future, and intangible* benefits in the form of additional product sales and reduced stock-out costs. The reduction in cycle stock and pipeline inventory are more tangible cost savings. While this effect would be interesting to explore, we do not believe that most game participants assessed inventory savings any differently that they did other cost savings.

Appendix

The BPI and the Votes Needed to Win an Election

The classic BPI assumes the number of votes needed to win the election to be the same for either choice. This does not appear to be rich enough to model the OsG, which is asymmetric in that COMMON requires 51 or more of the 100 votes but UNIQUE can prevail with just 50. Such a case arises because the total number of votes is even, and the tie-vote scenario is awarded preferentially to one of the choices. In fact, for an even number of votes (call it n) where a simple majority wins and abstentions are prohibited, the BPI values for a winning threshold of $n/2$ will be same as for $(n/2)+1$. This is because a defection that turns a winning coalition (for COMMON, in the OsG) into a losing one will simultaneously turn the complementary losing coalition into a winning one (for UNIQUE). When viewed in reverse, this maps one-to-one to the conversion of each winning UNIQUE coalition into a losing one. So, while the designation of which option will win in case of a tie will certainly change the likelihoods of winning, it will not change how power is distributed among the individual blocks in the Banzhaf sense. Thus, we may calculate the BPI for the OsG in the standard way with either 51 or 50 votes as the requirement for the win.

References

- Amaral, J., C. A. Billington, A. A. Tsay. 2006. Safeguarding the promise of production outsourcing. *Interfaces* 36(3): 220–233.
- Amaral, J., A. A. Tsay. 2008. The Outsourcing Game: A teaching simulation of power dynamics in outsourced supply chains. *Operat. Manag. Educ. Rev.* 2(1): 1–27.
- Banzhaf, J. F. I. 1965. Weighted voting doesn't work: A mathematical analysis. *Rutgers Law Rev.* 19(2): 317–343.
- Banzhaf, J. F. I. 1968. One man, 3,312 votes: A mathematical analysis of the electoral college. *Villanova Law Rev.* 13(2): 304–346.
- Bardhan, I., S. Mithas, S. Lin. 2007. Performance impacts of strategy, information technology applications, and business process outsourcing in U.S. manufacturing plants. *Prod. Operat. Manag.* 16(6): 747–762.
- Barthelemy, J. 2003. The seven deadly sins of outsourcing. *Acad. Manag. Exec.* 17(2): 87–100.
- Bazerman, M. H., J. R. Curhan, D. A. Moore, K. L. Valley. 2000. Negotiation. *Ann. Rev. Psychol.* 51(1): 279–314.
- Bendoly, E., K. L. Donohue, K. L. Schultz. 2006. Behavior in operations management: Assessing recent findings and revisiting old assumptions. *J. Operat. Manag.* 24(6): 737–752.
- Bowen, L. 2001. Power and weighted voting systems. Available at <http://www.Ctl.Ua.Edu/math103/power/wtvoting.Htm> (accessed date May 25, 2008).
- Camerer, C., R. H. Thaler. 1995. Anomalies: Ultimatums, dictators and manners. *J. Econom. Persp.* 9(2): 209–219.
- Carnegie, D. 1981. original copyright 1936. *How to Win Friends and Influence People*. Revised Edition. Pocket Books, New York.
- Cialdini, R. 2007. *Influence: The Psychology of Persuasion*. First Collins Business Essentials Edition. HarperCollins, New York.
- Dawkins, R. 2006. *The Selfish Gene*. Oxford University Press, New York.
- Doig, S. J., R. C. Ritter, K. Speckhals, D. Woolson. 2001. Has outsourcing gone too far? *McKinsey Q.* 26(4): 25–37.
- Dubey, P., L. S. Shapley. 1979. Mathematical properties of the Banzhaf Power Index. *Math. Operat. Res.* 4(2): 99–131.
- Earl, M. J. 1996. The risks of outsourcing IT. *Sloan Manag. Rev.* 37(3): 26–32.
- Fehr, E., S. Gächter. 2000. Fairness and retaliation: The economics of reciprocity. *J. Econom. Persp.* 14(3): 159–181.
- Fehr, E., S. Gächter. 2002. Altruistic punishment in humans. *Nature* 415: 137–140.
- Gelman, A., J. N. Katz, F. Tuerlinckx. 2002. The mathematics and statistics of voting power. *Stat. Sci.* 17(4): 420–435.
- Gilovich, T., D. Griffin, D. Kahneman. 2002. *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge University Press, New York.
- Kagel, J. H., A. E. Roth. 1995. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ.
- Kahneman, D. 2003. Maps of bounded rationality: Psychology for behavioral economics. *Am. Econom. Rev.* 93(5): 1449–1475.
- Kahneman, D., J. L. Knetsch, R. H. Thaler. 1991. Anomalies: The endowment effect, loss aversion, and status quo bias. *J. Econom. Perspect.* 5(1): 193–206.
- Lakenan, B., D. Boyd, E. Frey. 2001. "Why Cisco fell: Outsourcing and its perils." *Strategy+Business* (Q3 2001).
- Lax, D., J. K. Sebenius. 1986. *The Manager as Negotiator: Bargaining for Cooperation and Competitive Gain*. Free Press, New York.
- Lax, D., J. K. Sebenius. 2006. *3-D Negotiation: Powerful Tools to Change the Game in Your Most Important Deals*. Harvard Business School Press, Boston, MA.
- Lemisch, L. J. 1961. *Benjamin Franklin: The Autobiography and Other Writings*. Signet Classic, New York.
- Lieberman, V., S. M. Samuels, L. Ross. 2004. The name of the game: Predictive power of reputations versus situational labels in determining Prisoner's Dilemma game moves. *Personality Soc. Psychol. Bull.* 30(9): 1175–1185.
- Livingston, M. 2003. Banzhaf Power Index. Available at <http://www.Cs.Unc.Edu/~livingst/banzhaf/> (accessed date July 18, 2006).
- Machiavelli, N. 1515. *The Prince* (translated by W.K. Marriott). Available at <http://www.gutenberg.org/files/1232/1232-h/1232-h.htm> (accessed date February 12, 2009).
- Marx, L. M., G. Shaffer. 2001. Bargaining power in sequential contracting. Simon School of Business working paper No. FR 01-09.
- Nagarajan, M., Y. Bassok. 2008. A bargaining framework in supply chains: The assembly problem. *Manag. Sci.* 54(8): 1482–1496.
- Nagarajan, M., G. Sosis. 2008. Game-theoretic analysis of cooperation among supply chain agents: Review and extensions. *Eur. J. Operat. Res.* 187(3): 719–745.
- Penrose, L. S. 1946. The elementary statistics of majority voting. *J. R. Stat. Soc.* 109(1): 53–57.
- Raiffa, H. 1982. *The Art and Science of Negotiation*. Harvard University Press, Cambridge, MA.
- Rapoport, A. 1970. *N-Person Game Theory: Concepts and Applications*. University of Michigan Press, Ann Arbor, MI.
- Roth, A. E. 1995. Bargaining experiments. Kagel, J. H., A. E. Roth eds. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ, 253–291.
- Sebenius, J. K. 1996. Sequencing to build coalitions: With whom should I talk first? Zeckhauser, R. J., R. L. Keeney, J. K. Sebenius eds. *Wise Choices: Decisions, Games, and Negotiations*. Harvard Business School Press, Boston, MA, 324–348.
- Sebenius, J. K. 2004. Mapping backward: Negotiating in the right sequence. *Negotiation* 7(6): 3–5.

- Shapley, L. S., M. Shubik. 1954. A method for evaluating the distribution of power in a committee system. *Am. Polit. Sci. Rev.* 48(3): 787–792.
- Susskind, L. “Making the most of multiparty negotiations,” *HBS Working Knowledge* (February 9, 2004).
- Thompson, L. 2008. *The Truth About Negotiations*. FT Press, Upper Saddle River, NJ.
- Thurm, S. 2007 “Behind outsourcing: Promise and pitfalls” *The Wall Street Journal*: B3.
- Tsay, A. A. 2009a. Business process outsourcing. Cochran, J. J. ed. *Encyclopedia of Operations Research and Management Science*. Wiley, Hoboken, NJ. (in press).
- Tsay, A. A. 2009b. Supply chain outsourcing. Cochran, J. J. ed. *Encyclopedia of Operations Research and Management Science*. Wiley, Hoboken, NJ. (in press).
- Tversky, A., D. Kahneman. 1974. Judgments under uncertainty: Heuristics and biases. *Science* 185(4157): 1124–1131.