Supply Intelligence[1]

Sridhar Seshadri[2] and Eitan Zemel

Information, Operations, and Management Sciences Department

Stern School of Business

New York University, New York, NY 10012

sseshadr@stern.nyu.edu, ezemel@stern.nyu.edu

Abstract

We examine supply chain scenarios in which obtaining complete information about suppliers is costly. In such scenarios, there is a trade-off between the costs of obtaining information and the benefits that accrue to the owners of such information. Moreover, there are alternate mechanisms available for extracting such information.

In order to study the relative merits of the different ways information can be acquired, we construct a general model which builds on the models of Dasgupta and Spulber (1985) and Chen (2001), who use auctions to extract information. Our model is considerably more general in terms of its assumptions and information requirements. In particular, the model can handle general cost and revenue functions, can allow for multiple variables, and is not dependent on functional properties such as continuity or convexity. The model is also very sparse in terms of the information that is available to the different parties a priori.

We show that despite the restricted assumptions, the buyer can still guarantee significant profit levels for herself, while at the same time induce the channel to perform efficiently, i.e., at its (centralized) optimum. We also show how the buyer can use *audits,* performed after the transaction is implemented, in order to increase her share of the channel profits, while at the same time maintaining the incentives for channel coordination. We show that the well known buyback contract for the newsvendor problem can be viewed as a special case of audit based contracts. Finally, we examine the behavior of audit-based auctions when the audits are biased. This allows us to examine situation in which effort can reduce costs, but where the cost of effort cannot be documented. Naturally, the results apply equally well in scenarios where the roles of the supplier and buyer are reversed.

(Original version March 14, revised May 19 2003)

---

[1] We gratefully acknowledge the comments of the participants of the seminar at the Stern School of Business and Fangruo Chen of Columbia University.

[2] Sridhar Seshadri's research is partially supported by grants under DMI-0200406.

1. Introduction

Supply chain scenarios often involve multiple decision makers (agents), each possessing only part of the information necessary for finding the optimal economic decision in a given situation. For instance, a buyer of a given product may have to choose from amongst several possible suppliers, each facing his own private cost function. If the buyer knew these cost functions, in principle, she could easily determine the optimal quantity to purchase from each supplier by solving an optimization problem. In the absence of complete information, the buyer must settle for less profitable actions, which are best given her partial information, but are not optimal overall. Naturally, in such a setting, the buyer may wish to spend some cost and effort in trying to obtain some of the relevant, missing information from the other agents. Clearly, if she is to be successful in such efforts, she must take into consideration the incentives and objectives of these agents.

There is a vast literature in the fields of Operations Management and Economics that addresses different aspects of this question. We survey some of the relevant streams of this literature in the next section. In this paper we investigate alternative *forms* of information that could be exchanged, and in particular, the central role that *audits* can play in this regard. Specifically, we compare the advantages and limitations of gathering information via audits, that are performed *after* the transaction have occurred, to the standard approach utilized in many papers in this area that emphasizes the role of a priori information, that is learned by the agents before the contract is finalized. To that end, we develop an extremely general model of a two-stage supply chain. In setting up this model, we attempt to minimize, to the extent possible, the assumptions with regard to the cost and revenue functions faced by the agents, as well the assumptions regarding the level of *prior information that each agent has with respect to the other agents*. In particular, we do not assume any specific functional form of the cost or revenue functions, nor do we rely on functional properties such as continuity, convexity, differentiability, etc. In addition, we allow the supply chain transaction between the buyer and the supplier to involve any number of variables, such as quantity, quality, flexibility, delivery frequency, etc. In addition, we allow private variables such as supplier effort and retailer pricing discretion.

We study the buyer's profit in such a setup. Where possible, we compare it to the profit achieved in the classic model of Dasgupta and Spulber (1990) and Chen (2001). We then examine in particular the impact of *audits* as a potentially effective method for achieving coordination. We also examine the impact on the buyer's profit when the audit is incomplete (biased) due to factors such as effort.

2. Literature (supply chain contracts)

Our paper draws from four main streams of literature. The first stream is the rich and extensive literature on coordination and contracts in supply chain management. Excellent surveys of this literature can be found in the book by Tayur et al (1998), (for example Anupindi and Bassok 1998), as well as in Cachon (2003), Chen (2002) and Elmaghraby (2000). The second stream, denoted as "principal-agent" models, concerns models of incentives and asymmetric information between two agents. The book by Laffont and Tirole (1994) deals extensively with such models developed in the context of regulation and procurement. The third stream is the literature on auction theory and optimal mechanism design (Klemperer (1999), Maskin and Riley (1984), McMillan and McAfee (1984), Milgrom (1981), Myerson (1981), Riley and Samuelson (1981)). Finally, we rely in particular on the work of Dasgupta and Spulber (1990), who were one of the earliest to introduce the notion of an auction into a supply chain type model (their work was in the context of government contracting). This model was later further developed and refined by Chen (2001). We refer to the models of Dasgupta and Spulber (DS), and Chen collectively as DSC. Below we summarize the DSC model.

3. *The DSC Model*

In the DSC model there is a single buyer who wishes to procure a quantity Q from (at most) one of n firms. The buyer has a benefit function (expected revenue for example) $R(Q)$. The suppliers have cost functions given by $C_i(Q)$. Thus, if supplier i is chosen, and the quantity Q were purchased, the entire supply chain profit is equal to $R(Q)\text{-}C_i(Q)$. DSC assume the following:

(1) The cost function $C_i(Q)$ is known only to supplier i. The benefit function $R(Q)$ is known only to the buyer.

(2) The buyer designs the contract, i.e., she can make any "take it or leave it" offer to any of the suppliers, who will accept it if its value to him is at least zero.

(3) Buyer and suppliers are risk-neutral.

(4) The function $R(Q)$ is concave, increasing and twice differentiable.

(5) The cost functions are of the form $C_i(Q) = C(Q,q_i) = K(q_i) + \int_0^Q c(q,q_i)dq$, with a private parameter $q \in [\underline{q},\overline{q}]$. Chen assumes that the cost functions are linear, with fixed cost equal to zero. DS assume that $K(q)$ > 0 and that the marginal cost is increasing in Q, ($c_Q > 0$), as well as, increasing and convex in the type parameter, (c_q ³ 0, and c_{qq} ³ 0). Finally, DS also assume that the costs can be parameterized such that c_{qQ} ³ 0.

(6) The sellers are drawn independently and identically from a continuous distribution $F(q)$, where $q \in [\underline{q},\overline{q}]$ (i.e., the support of F is this closed interval). In addition, $H(q) = F(q)/F'(q)$, is increasing in q.

(7) The supplier knows the cost functions $C(Q,q)$ and $H(q)$.

Note that assumptions (4)-(6), although standard, may be quite restrictive. For instance, assumption (5) implies that the suppliers can be sorted from the cheapest to the most expensive, and that this ordering is independent of Q. This excludes situations in which different suppliers have cost advantages in different parts of the cost curve. Also, assumption (7) requires that considerable amount of information be available to the buyer in advance. Naturally, in practice, the buyer will need to spend considerable effort and cost obtaining this information. However, given these assumptions, DSC utilize very elegant reasoning to derive the optimal auction mechanism for the buyer.

To appreciate the results of DSC, note that it follows easily from their assumptions that if the types θ_i were known to the buyer (contrary to assumption), she could have chosen the best (cheapest) supplier and computed the optimal quantity Q. Presenting this supplier with a "take it or leave it" offer to produce the right quantity, she could have expropriated the entire profit of the channel. Absent this information, DSC show how auctions can be used as a mechanism for the buyer to obtain information in order to maximize her share of the channel profits. DS first analyze a model in which the buyer selects a fixed quantity Q, and then uses an auction to purchase this quantity at the lowest cost. In this case, the problem essentially reduces to that of buying a single object (to wit, a package of Q units) that has different value (cost) to each of the bidders (suppliers). Thus, the framework of private value auctions (Maskin and Riley(1984), Vickrey (1961)) is applicable. This however, begs the question of how the quantity Q is to be determined. Moreover, even if the optimal Q were known to the buyer in advance, DS show that she can do better by allowing the quantity purchased to depend on the type of the winning bidder. The elegant trick of DSC is to offer for sale via auction the ***right to determine Q.*** Relying extensively on assumptions (4)-(6), the optimal auction mechanism for the buyer is determined.

How well does the buyer fare under DSC? Examine the total channel profits. If supplier i were to be selected, and if he were to optimize Q with respect to his cost function, the channel profit would be:

$$\Pi_i = \max_Q \left(R(Q) - C(Q, q_i) \right).$$

Assume the suppliers are indexed in decreasing order of their profit Π, and let the buyer's optimal profit be denoted by Π^{*B}. Obviously, the maximum conceivable profit for the buyer is Π_1. Given the assumptions of DSC, however, this level of profit cannot be achieved accept for degenerate cases. However,

Proposition 1 (DS):

$$E\left(\Pi^{*B}\right) = n\int_{\underline{q}}^{q^{**}} (1 - F(q))^{n-1} F'(q)\left(R(Q) - C(Q,q) - C_q(Q,q)H(q)\right)dq$$

where

(3.1) $$R(Q(q^{**})) = C(Q(q^{**}),q^{**}) + C_q(Q(q^{**}),q^{**})H(q^{**})$$

and

(3.2) $$R'(Q(q)) = C_Q(Q,q) + C_{qQ}(Q(q),q)H(q), q \in \left[\underline{q},\bar{q}\right],$$

where q^{**} is chosen according to (3.1) in order to exclude suppliers with higher cost parameters (q^{**} is set equal to $\bar{q}$ if no such θ exists) and (3.2) gives the quantity that is produced if the winning supplier has cost parameter θ.

The second term on the right hand side of (3.2) is positive by assumptions (4) - (6). Thus, the quantity contracted is less than optimal (*R* is concave). This yields, $\Pi^{*B} \le \Pi_1$. On the other hand, it can be shown that the difference between the highest and second highest value, given θ_1, is exactly equal to $C_q(Q,q_1)H(q_1)$ (see Chen). Further, according to (3.2), the buyer's profit is maximized given that the winning supplier has cost parameter equal to q. In other words, her profit is being maximized conditional to the winning supplier's cost parameter. Therefore, under the optimal DSC mechanism the expected value of the buyer's profit will be higher than if we were to maximize the unconditional expected value, which equals Π_2. Thus, under assumptions (1)-(7), $\Pi_2 \le \Pi^{*B} \le \Pi_1$.

The mechanism for maximizing the buyer's profits relies crucially on assumptions (4) – (6), and on the information available to the buyer per assumption (7). We also note that under the optimal mechanism, the channel is *not efficient*, i.e., the total channel profit is less than Π_1. In the next section, we examine what could be achieved when assumptions (4)-(7) do not hold in their entirety.

4. *The General Model*

We now present a model that generalizes the model considered by DSC. We assume a single buyer who wishes to procure a "package" (of goods) from (at most) one of n firms. However, while DSC specify the package in terms of a single (and continuous) variable, namely the quantity Q, we allow the interaction between supplier and buyer to depend on a multitude (k) of parameters. These parameters could relate to quantity, quality, capacity, delivery time, delivery frequency, etc. We denote this collection of parameters by the vector, $\mathbf{x} = (x_1, x_2, \ldots x_k)$. We refer to $\mathbf{x}$ as the ***contract***. Let X be the set of all possible contracts. Some or all of the components of $\mathbf{x}$ could be discrete.

Given a contract $\mathbf{x} \in X$, we assume that the buyer's revenue (benefit) can be expressed as $R(\mathbf{x})$. Similarly, we denote the *cost* for delivering contract $\mathbf{x}$ by supplier i as $C_i(\mathbf{x})$. For each supplier, and each level of the contract $\mathbf{x}$, denote the total channel profit as $P_i^C(\mathbf{x}) = R(\mathbf{x}) - C_i(\mathbf{x})$.

We make the following assumptions:

(1*) The cost function $C_i(\mathbf{x})$ is known only to supplier i. The benefit function $R(\mathbf{x})$ is known only to the buyer.

(2*) The buyer designs the contract, i.e., she can make any "take it or leave it" offer to any of the suppliers, who will accept it if its value to him is at least zero.

(3*) Buyer and suppliers are risk-neutral.

(4*) For every supplier, $\mathbf{x}_i^* = \operatorname{argmax}\{P_i^C(\mathbf{x}) : \mathbf{x} \in X\}$ exists and $P_i^C(\mathbf{x}_i^*)$ is finite. The maximizing argument $\mathbf{x}_i^*$ need not be unique.

Note that assumptions (1*)-(3*) are the exact analog of assumptions (1)-(3) of DSC, with the more general argument $\mathbf{x}$ replacing Q. However, assumption (4*) is a considerable relaxation of assumptions (4)-(6): we are not imposing any functional restrictions on $R(\mathbf{x}), C_i(\mathbf{x})$, (such as convexity, differentiability etc.), or on the structure of the set X. In particular, the components of $\mathbf{x}$ could be discrete or continuous. Additionally, we do not require that the buyer have any information on the cost functions faced by the supplier. We refer to assumptions (1*)-(4*) as *the reduced* set of assumptions, and to assumptions (1)-(8) as the DSC assumptions.

Before we examine what is achievable under assumptions (1*)-(4*), it is useful to elaborate about the generality of the revenue and cost functions used in our model. In many models, these functions are assumed to be the primitive elements of the model. For example, in numerous papers the cost functions are assumed to be linear. However, in some situations, we may wish to include, in addition to the variables **x** which define the contract between the supplier and the buyer, other essential variables, that are available and known to only one of the parties. In addition, we may want to include in the analysis the effect of timing and random events. Since we make almost no assumptions on the cost and revenue functions, the framework we present can accommodate such variations easily, by letting $R(\mathbf{x})$ and $C(\mathbf{x})$ depend upon more elementary elements of the model. For example, if the effects of *randomness* were critical to the logic of the model, we let

$$R(\mathbf{x}) = E[R(\mathbf{x}, D)]$$

where the expectation is with respect to D, a random variable with a given distribution. As an example, in the case of the newsvendor model the revenue function is of the form

$$R(Q) = p \bullet E[Min(Q, D)]$$

where the one-dimensional variable Q represents the order quantity, and p is the unit selling price. Additionally we may want to include in the model private variables, controlled exclusively by one of the parties (such as pricing and discounting decisions, that are determined by the buyer). Those could be modeled by setting

$$R(\mathbf{x}) = Max_v[R(\mathbf{x}, v)]$$

where v is a vector of private variables, controlled by the buyer. For instance, if the selling price p is a decision variable for the buyer, and if $Q(p)$ is the quantity sold at price p, then for a given order quantity x, the revenue can be written as

$$R(x) = Max_p[p * Q(p) : Q(p) \leq x].$$

Similarly, in the case of a supplier who faces various production technologies characterized, say, by a tradeoff between fixed costs and variable unit costs, *C(x)* could represent the minimal cost, over the various available technologies, of producing a given quantity *x*. Obviously, the number of modeling options that could be accommodated is quite large for both $R(\mathbf{x})$ and $C(\mathbf{x})$, as the framework allows for any combination of the Max{ } and E{ } operators.

How well can the buyer do under the reduced set of assumptions? While this question is open, we can examine the predictions of cooperative game theory. Specifically, consider the cooperative game V where the value V(S) of each coalition S is equal to the maximal channel profits that could be generated by the members of S. The core of the game V is an allocation of the total channel profits to each agent (buyer and suppliers) such that each coalition (subset of agents) gets at least as much as they could generate on their own (if such allocation exists). Thus, the *core* of a cooperative game does not take into account the effects of asymmetric information, but rather is focused on the incentives of the various coalitions of players. The following proposition is from Seshadri and Zemel (2003). We denote the surplus allocated to the buyer by y^B, and to seller *i* by y_i. We assume that the player are indexed in decreasing order of the profits Π_i:

Proposition 2: Given the reduced set of assumptions (1), (2*), (3) and (4), the core of V is not empty. Further, the core is the closed interval

$$\begin{aligned} &y^B \in [\Pi_2, \Pi_1] \\ &y_1 = \Pi_1 - y^B \quad . \\ &y_i = 0, i \neq 1 \end{aligned}$$

5. Auctions

Recall that $\Pi_i = Max\{P_i^C(\mathbf{x}) : \mathbf{x} \in X\}$, and that under assumptions (1)-(7), the buyer's profit Π^B satisfies the relation $\Pi_2 \leq \Pi^B \leq \Pi^{*B} \leq \Pi_1$, where Π^{*B} is the buyer's optimal profit of Proposition 1. In this section we examine the fraction of channel profits that the buyer can extract using less information than is required by DSC. The following simple proposition asserts that the buyer can achieve the lower bound of the interval $[\Pi_2, \Pi_1]$ without any difficulty:

Proposition 3: Given the reduced set of assumptions (1*)-(4*) the buyer can achieve the profit level Π_2. The entire channel profits will be Π_1, i.e., the channel will be efficient (coordinated).

Proof: The buyer publishes a price schedule $R(\mathbf{x})$. She then offers for sale a *commitment to buy from the winning supplier any feasible* $\mathbf{x} \in X$ *for a price* $R(\mathbf{x})$. Thus, the winning supplier owns the option to sell any value $\boldsymbol{x}$ and get paid according to the price schedule $R(\mathbf{x})$. The winning supplier is the one with the highest bid b_1 and the payment to the buyer is the amount bid by the second highest bidder, b_2.

The value of winning such a contract for each supplier is Π_i, since once the contract is won, it is in his best interest to select the contract level $\mathbf{x}$ that maximizes his profit. Also, given that the auction is of the "second price" type, it follows that it is a dominant strategy for each buyer to bid exactly his value, i.e., $b_i = \Pi_i$ for every supplier. Thus, the winning supplier's profit is $\Pi_1 - \Pi_2$ and the buyer's profit is Π_2. □

Recall that under the assumptions of DSC, the buyer can achieve Π^{*B}. Since $\Pi_2 \leq \Pi^{*B}$, it is clear that DSC allow the buyer to achieve higher profits than those of the reduced set. Of course, the buyer of DSC has been "given" more information (assumption (7)), and the model is much more restrictive (assumptions (4)-(6)). Also, note that while in Proposition 2 the first best outcome is achieved, i.e., the channel is efficient, this is not the case in Proposition 1. Thus, to obtain a level of profits higher than Π_2 under DSC, the buyer causes the total channel profits to decrease. Finally, notice that when the number of bidders is large, the interval between the first and second values, $\Pi_1 - \Pi_2$, is typically small.

6. Audits

We define an *audit* to be an after-the-fact, verifiable measure of the actual cost incurred by the supplier. For example, a summary of expenses could be submitted by the supplier and verified by an independent and trustworthy accountant. We denote the audited cost by $\widetilde{C}$. We do not require that the audit be precise - it may differ from the actual cost by a random error term[3]. However, we require that it to be unbiased, i.e., that the expected value of the error term is zero. Below we show that using an unbiased audit, the buyer can obtain any value of profit in the semi-open interval $[\Pi_2, \Pi_1)$. The performance of biased audits, (for instance in cases where the supplier can exert effort to reduce the observed cost, but the cost of the effort is not captured in the audit), is studied subsequently.

Note that the cost of performing the audit is not taken into account in the following proposition. We examine the costs of audits, as well as the costs of estimating the cost function as per assumption (7) of DSC, later in this section

Proposition 4: Given assumptions (1*)- (4*), and the existence of an unbiased audit, the buyer can achieve any profit in the interval $[\Pi_2, \Pi_1)$. The payoff to the winning supplier is $\Pi_1 - \Pi^B$, i.e., the channel is efficient.

Proof: We use the theory of auctions with *contingent payments* introduced in Hansen (1985) to establish the proposition. The buyer publishes a price schedule $R(\mathbf{x})$ and selects a value $m \in (0,1]$. He then offers for sale a commitment to buy from the winning supplier any feasible $\mathbf{x} \in X$ for a price $R(\mathbf{x})$. The winning

[3] The randomness could be in the cost, in the audit, or in both.

supplier is the one with the highest bid b_1. He pays the buyer the amount $m \bullet b_2 + (1-m) \bullet \left[R(\mathbf{x}) - \widetilde{C} \right]$.

Note that the payment by the winning supplier is independent of the winning supplier's bid. Thus, it is a dominant strategy for each supplier to bid exactly his value, i.e., $b_i = \Pi_i$ for every supplier. Also, once the contract is won, it is in his best interest to select the contract level $\mathbf{x}$ which maximizes his profit since his objective function amounts to $m \bullet [R(\mathbf{x}) - C_1(\mathbf{x})]$. This is since for any $\mathbf{x}$ chosen by the winning supplier, the expected value of $\widetilde{C}$ is equal to $C_1(\mathbf{x})$. Thus, the winning bidder's profit is $m \bullet [\Pi_1 - \Pi_2]$ and the buyer's profit is $m \bullet \Pi_2 + (1-m) \bullet \Pi_1$. By varying m in the interval (0,1], the buyer's profits can cover the entire interval $[\Pi_2, \Pi_1)$. □

Remark: It does not benefit the buyer to turn away any bidder for whom $\sup_{\boldsymbol{x}} \boldsymbol{R}(\mathbf{x}) - C_i(\mathbf{x})$ is non-negative. In particular, if *for each supplier* there is an option that yields zero profit ($\exists \mathbf{x}_{i0} : R(\mathbf{x}_{i0}) - C_i(\mathbf{x}_{i0}) = 0$), then no bidder is turned away. This is unlike the DSC mechanism in which the upper cut-off value θ^{**} might exclude suppliers that are profitable to do business with by themselves.

Why do audits enable the buyer to obtain profits that are higher than the maximum Π^{*B} of Proposition 1? The key is her ability to make the supplier's payoff contingent upon the revealed costs (information). This idea has been discussed by Townsend (1979), Baron and Besanko (1984), as well as in Laffont and Tirole (1994).

In the case of a single buyer and a single supplier such as in the case of a regulator and the regulated firm, Baron and Besanko point out that the buyer (regulator) typically does not possess "full information" about the supplier (regulated firm), and that welfare can be improved through audit. (Baron and Besanko's analysis is from a welfare maximization perspective where the objective is to maximize the weighted sum of consumer and producer surpluses.)

The theory to analyze the case of a single buyer (or seller) transacting with many sellers (or buyers) was proposed in Milgrom (1985) and elaborated upon in Samuelson's (1987) response to Hansen (1985). Hansen (also alluded to in Proposition 4) modeled a "target" firm that faces takeover bids from *n* acquiring firms. Hansen shows that instead of simply selecting amongst bids to acquire it, the target firm stands to profit more if the acquiring firms bid the *share* of the merged firm that they are willing to give to the target firm. Thus, the target firm makes the price contingent on the value of the merger to the acquiring firm. Samuelson comments that it is the "linkage" principle, whereby the price is linked to a variable that is affiliated with the private information of the bidder, that reduces the bidder's rent (enjoyed by bidder possessing private information). This linkage principle is seen to provide higher expected profit in Proposition 4 compared to the optimal auction mechanism (without the linkage) of Proposition 1.

Naturally, Proposition 4 could be used in situations when the roles of the supplier and buyer are reversed, i.e., there is one supplier and several buyers, and the supplier proposes the contract as well as audits the buyer's revenue. In that case the supplier offers to sell any $\mathbf{x}$ for $C(\mathbf{x})$, and the winning buyer pays $m \bullet b_2 + (1-m) \bullet \left[\tilde{R} - C(\mathbf{x}) \right]$, where $\tilde{R}$ is the unbiased estimate of the revenue revealed by the audit.

As an application of the use of Proposition 4 in this manner, consider the ***buyback contract*** (Pasternack (1985), Cachon 2003) for the newsvendor model, one of the most cited contracts in the supply chain coordination literature. To see how it could be obtained from Proposition 4, note first that Proposition 4 holds even in the case of a single supplier. In that case, the contract is awarded without a bid, since the question of selecting the best supplier is moot. Rather, the purpose of the mechanism is to induce the buyer to purchase the optimal level of *x*. Note that b_2 in that case is taken to be zero. Second, note that under the assumptions of Pasternack, $C(Q) = cQ$, and $\tilde{R} = p \bullet E(Q - \text{Ret})$ where *p* is the selling price (known to the supplier), Ret is the number of units not sold (the "returns" under the buyback policy), and the operator E signifies the expected value, with respect to Ret. Thus knowing the actual Ret in this case provides a revenue "audit". The payment to the buyer under Proposition 4 is

$$C \bullet Q + (1-m) \bullet [\tilde{R} - C \bullet Q] = m \bullet C \bullet Q + (1-m) \bullet p \bullet [Q - \mathrm{Re}t]$$
$$= [m \bullet C + (1-m) \bullet p] \bullet Q - (1-m) \bullet p \bullet \mathrm{Re}t$$

which is precisely the payment scheme under buybacks.

6.1 Cost of Information

As noted earlier in the section, the costs of *auditing* the winning supplier is not included in the derivation of the results of proposition 4, just as the cost of *estimating* the functions $C(Q,q)$ and $H(q)$ is not taken into account in the analysis of DSC. Naturally, in either case, if these costs are high relative to the profits involved, the approach can not be used. One would expect that the relative costs of auditing and estimating are highly dependent on the situation analyzed. We note, however, a few differences between the characteristics of these two approaches to obtaining information:

(1) Auditing involve only the winning supplier. Thus, the cost of audit is independent on the number of bidders. In contrast, the estimation of the cost functions involves obtaining information concerning all the bidders. Thus, we can expect auditing to be relatively more useful in situations involving more bidders.

(2) Audits involve the actual cost, at a given level of output, while estimation involve the entire cost *function*. Thus, we can expect auditing to be relatively more useful in situations involving complex cost functions.

(3) Without any information, the buyer can obtain profit level Π_2 as per Proposition 3. The maximal profit with auditing is Π_1. Thus, if the costs of auditing is large relative to $\Pi_1 - \Pi_2$, it is not economical to audit.

(4) The maximal profit under DSC is Π^{B*}. Thus, if the costs of estimating is large relative to $\Pi^{B*} - \Pi_2$, it is not economical to estimate.

(5) In some situations, especially if penalties could be imposed, audits could be replaced by self reporting supplemented by random audits. If the penalty could be set high, the frequency of audit could be set low, reducing the actual costs of audits.

6.2 An example

In order to compare the results of Propositions 1, 2 and 4, consider a slight modification of the newsvendor example analyzed in Chen 2001, 2002. Demand is uniformly distributed in the interval [0,1]. The product sells for \$2, and unsold product has zero salvage value . The buyer orders an amount *Q*. Thus, the (expected) revenue function is given by $R(Q) = 2Q-Q^2$. The cost function is linear, viz., $C(Q,\theta) = \theta Q$, where the cost parameter θ is distributed uniformly in the interval [0.1,1.1] so that $H(q) = F(q)/f(q) = q-0.1$. There are *n* firms.

Applying equations (3.1) and (3.2) to this model we obtain the optimal Q for each *q*, as well as the upper cutoff value when the buyer uses the Auction without Audit (or AWA) mechanism:

$$Q^{AWA}(q) = 1.05 - q$$
$$q^{**} = 1.05.$$

Also, the expected profit for the buyer equals:

$$p^{*B} = n\int_{0.1}^{1.05}(1.1-q)^{n-1}(2Q(q)-Q(q)^2 - C(Q(q),q) - C_q(Q(q),q)H(q))dq$$
$$= n\int_{0.1}^{1.05}(1.1-q)^{n-1}(1.05-q)^2 dq$$
$$= 0.95^2 - \frac{1.9}{(n+1)} + \frac{2(1-0.05^{n+2})}{(n+1)(n+2)}.$$

On the other hand, Auction and Audit (A&A) will yield the first best solution, namely, the quantity will be given by the profit maximizing value:

$$Q^{A\&A}(q) = \text{argmin}_Q \{2Q - Q^2 - q\,Q\} = 1 - q/2.$$

The cut-off value in this case is $q^{**} = 1.1$. (In particular notice that when using AWA the optimal quantity will deviate the most from the ex post optimal quantity under A&A when the q of the winning supplier is high -the two are equal only when $\theta = 0.1$.) The maximum expected profit for A&A will be

$$p_1 = n\int_{0.1}^{1.1}(1.1-q)^{n-1}(2Q(q)-Q(q)^2-C(Q(q),q))dq$$

$$= n\int_{0.1}^{1.1}(1.1-q)^{n-1}(1-q/2)^2 dq = 0.95^2 - \frac{0.95}{n+1} + \frac{0.5}{(n+1)(n+2)}.$$

Let $\Pi_i(n)$ stand for the i-th highest value when there are n suppliers. Then, the value of $\Pi_2(n) = n\,\Pi_1(n-1) - (n-1)\,\Pi_1(n)$. Values of Π_1, Π^{*B} and Π_2 are listed in Table 1 for various values of n and also shown in Figure 1.

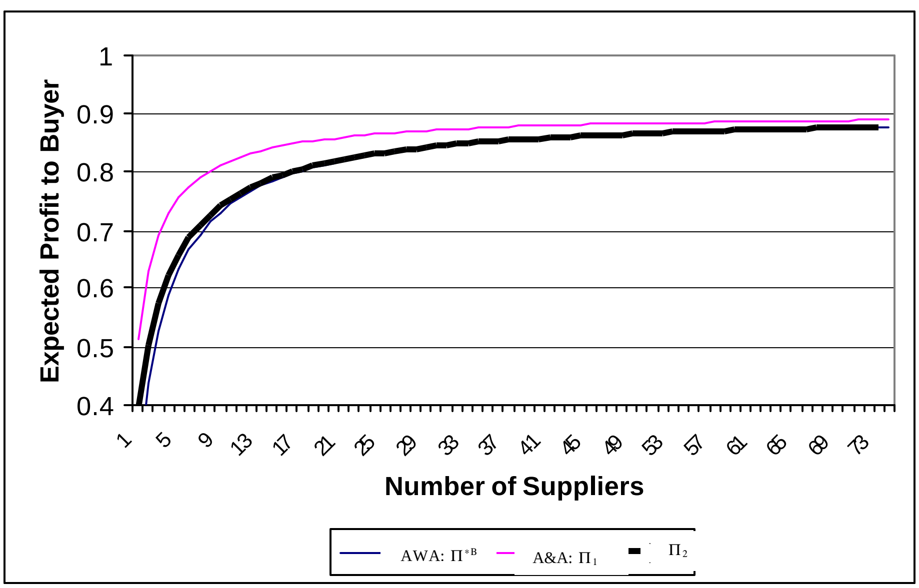

Figure 1: Expected profit to the buyer under AWA and A&A

n	$E[\Pi_2]$	AWA: $E[\Pi^{*B}]$	A&A: $E[\Pi_1]$	% inc. (A&A/ over AWA
2	0.3942	0.4358	0.6275	43.98%
5	0.6215	0.6335	0.7561	19.36%
10	0.7411	0.7449	0.8199	10.07%
20	0.8153	0.8164	0.8583	5.14%
30	0.8427	0.8432	0.8724	3.45%
40	0.8570	0.8573	0.8796	2.60%
50	0.8658	0.8660	0.8841	2.09%

Table 1: Selected values of expected profit to buyer under AWA and A&A

Consider the impact of the number of suppliers, *n*, on the buyer's expected profit. Clearly, when *n* is small, audit is quite advantageous since the distance between Π^{*B} and Π_1 is quite large. For large values of *n*, the distance between Π_1 and Π_2 is rather small so the three approaches (AWA, A&A, and Proposition 2) give similar results (asymptotically they approach 0.95^2). Consider now the cost of gathering information. In the case of Π^{*B}, the cost imposed by assumptions (1)-(8) will increase with *n*. In contrast, the cost of auditing remains unchanged.

7. *Biased Audits and the Effects of Effort*

In this section, we consider the situation in which the audit is *biased* because the auditor is unable to observe (and thus reimburse for) certain costs. We make the following two assumptions:

9) We extend the general model to accommodate this by writing

$$C_i(\mathbf{x},e) = D_i(\mathbf{x},e) + E_i(\mathbf{x},e),$$

where, $\boldsymbol{e}$ is a collection of unobservable variables that affect the cost, $D_i(\mathbf{x},e)$ is the observable part of the cost, and $E_i(\mathbf{x},e)$ (≥ 0) is the unobservable part. Moreover, we can conveniently set $E_i(\mathbf{x},0) = 0$. We refer to e as *effort*, and to $E_i(\mathbf{x}, e)$ as the *cost of effort.*

10) We assume that the auditor obtains an "unbiased" estimate of $D_i(\mathbf{x},e)$, that is denoted as $\widetilde{D}_i$, but that the cost of effort cannot be audited.

How will the mechanism of Proposition 4 fare, if the audit is biased? Specifically, the buyer publishes a price schedule $R(\mathbf{x})$. She then offers for sale a commitment to buy from the winning supplier any feasible $\mathbf{x} \in X$ for a price $R(\mathbf{x})$. Supplier i submits bids b_i. The winning supplier is the one with the highest bid: b_1. He pays the buyer the amount $m \bullet b_2 + (1-m) \bullet \left[R(\mathbf{x}) - \widetilde{D}_1 \right]$. Let

$$d_i(m) = \max_{\mathbf{x},e}\left(R(\mathbf{x}) - C_i(\mathbf{x},e) - \frac{1-m}{m} E_i(\mathbf{x},e) \right), \mu \in (0,1].$$

Let the suppliers be ordered in the descending order of the d_i's (this ordering might depend on μ). Let $\mathbf{x}^*(\mu)$ and $e^*(\mu)$ denote the maximizing values for the supplier with the largest value of b.

Proposition 5: Given assumptions (1), (2*), (3), (4), (9), (10), for every $m \in$ ***(0,1]*** the buyer's profits are given by:

$$p^B(m) = m \bullet d_2(m) + (1-m) \bullet d_1(m) + \frac{1-m}{m} E_1(x^*(m), e^*(m)). \tag{7.1}$$

Proof: We drop the dependence on μ for ease of notation. Consider supplier i's problem. If he wins the auction, his expected payoff (the expectation is with respect to the uncertainty in the audited cost) is given by:

$$-m \bullet b_2 - (1-m) \bullet E\left[R(\mathbf{x}) - \widetilde{D}_1 \right] + R(\mathbf{x}) - C_1(\mathbf{x},e)$$

$$= -m \bullet b_2 + m \bullet \left[R(\mathbf{x}) - C_1(\mathbf{x},e) - \frac{1-m}{m} E_1(\mathbf{x},e) \right]$$

$$= m[(R(\mathbf{x}) - C_1(\mathbf{x},e) - \frac{1-m}{m} E_1(\mathbf{x},e)) - b_2]. \tag{7.2}$$

Thus, supplier i's dominant strategy is to bid

$$b_i = d_i = \max_{\mathbf{x},e}\left(R(\mathbf{x}) - C_i(\mathbf{x},e) - \frac{1-m}{m} E_i(\mathbf{x},e) \right). \tag{7.3}$$

Let $\mathbf{x}^*$ and $\boldsymbol{e}^*$ denote the maximizing values for the winning supplier.

The buyer's expected profit is equal to

$$
\begin{aligned}
& m \bullet d_2 + (1-m) \bullet \left[R(\mathbf{x}^*) - D_1(\mathbf{x}^*, e^*) \right] \\
&= m \bullet d_2 + (1-m) \bullet \left[R(\mathbf{x}^*) - C_1(\mathbf{x}^*, e^*) \right] + (1-m) E_1(\mathbf{x}^*, e^*) \\
&= m \bullet d_2 + (1-m) \bullet d_1 + \frac{1-m}{m} E_1(\mathbf{x}^*, e^*).
\end{aligned}
$$

□

Note that the sum of the buyer's and the supplier's expected profits equals:

$$
\begin{aligned}
m \bullet d_2 + (1-m) \bullet d_1 + \frac{1-m}{m} E_1(\mathbf{x}^*, e^*) + m \bullet (d_1 - d_2) &= d_1 + \frac{1-m}{m} E_1(\mathbf{x}^*, e^*) \\
&= R(\mathbf{x}^*) - C_1(\mathbf{x}^*, e^*).
\end{aligned}
$$

Thus, the outcome is typically not efficient; the failure to observe the cost of effort induces the channel to exert sub-optimal level of effort.

Proposition 6:

(a) The function $d_i(m)$ is increasing in m.

(b) Further, $d_i(1) = P_1$.

Proof: From the proof to Proposition 5,

$$
d_i = \max_{\mathbf{x},e} \left(R(\mathbf{x}) - C_i(\mathbf{x}, e) - \frac{1-m}{m} E_i(\mathbf{x}, e) \right).
$$

Due to the fact that $(1-m)/m$ is a decreasing function of m, for every fixed $\boldsymbol{x}$ and e and $\mu_1 < \mu_2$

$$
\left(R(\mathbf{x}) - C_i(\mathbf{x}, e) - \frac{1-m_2}{m_2} E_i(\mathbf{x}, e) \right) \geq \left(R(\mathbf{x}) - C_i(\mathbf{x}, e) - \frac{1-m_1}{m_1} E_i(\mathbf{x}, e) \right).
$$

Thus, the maximum of the left hand side of the inequality (over $\boldsymbol{x}$ and $\boldsymbol{e}$) is not smaller than the maximum of the right hand side. This proves (a). The proof of (b) follows from observing that when $m = 1$ and the definition of $d_i(m)$, the supplier maximizes the channel profit, $R(\mathbf{x}) - C_i(\mathbf{x}, e)$. □

The tradeoffs inherent in biased audits can now be articulated. When m is small, the buyer places heavy emphasis on the audit. In this case, the buyer's share approaches $d_1(0)$, i.e., she appropriates all the profits. In turn, the winning supplier obtains no profit and exerts very little effort. On the other hand, when m is close to one, the buyer places little emphasis on the audit. In this case the buyer's share approaches $b_2(1) = \Pi_2$ and the winning supplier receives $\Pi_1 - \Pi_2$ and exerts full effort. Thus the buyer gets a smaller share, but of a bigger pie. The behavior of the buyer profits for intermediate values of m can not be determined based on our reduced set of assumptions. We return to this issue in the next section. However, we can characterize the behavior of the channel profit with respect to m as follows:

Lemma 1: If the functions, $R(\mathbf{x})$, $C_i(\mathbf{x}, e)$, and $E_i(\mathbf{x}, e)$ are differentiable in their arguments, have a unique minimum, and the optimal value of $E_i(\mathbf{x}, e)$ is increasing in μ, then the channel profit is non-decreasing in μ.

Proof: Let $\mathbf{x}_i(m)$ and $e_i(m)$ be the optimal values for supplier i. Then,

$$
\begin{aligned}
& d\left(R(\mathbf{x}_i(m)) - C_i(\mathbf{x}_i(m), e_i(m))\right)/dm \\
&= d\left(R(\mathbf{x}_i(m)) - C_i(\mathbf{x}_i(m), e_i(m)) - \frac{1-m}{m} E_i(\mathbf{x}_i(m), e_i(m))\right)/dm + d\left(\frac{1-m}{m} E_i(\mathbf{x}_i(m), e_i(m))\right)/dm \\
&= -\frac{d}{dm}\left(\frac{1-m}{m}\right) \times E_i(\mathbf{x}_i(m), e_i(m)) + d\left(\frac{1-m}{m} E_i(\mathbf{x}_i(m), e_i(m))\right)/dm \\
&= \left(\frac{1-m}{m}\right)\frac{d}{dm} E_i(\mathbf{x}_i(m), e_i(m)) \geq 0.
\end{aligned}
$$

The second equality follows from the envelope theorem. The final inequality follows from the assumption of the proposition. □

A typical example of the relevant functions is shown in Figure 2 below. In the case depicted the optimal value of m is 1:

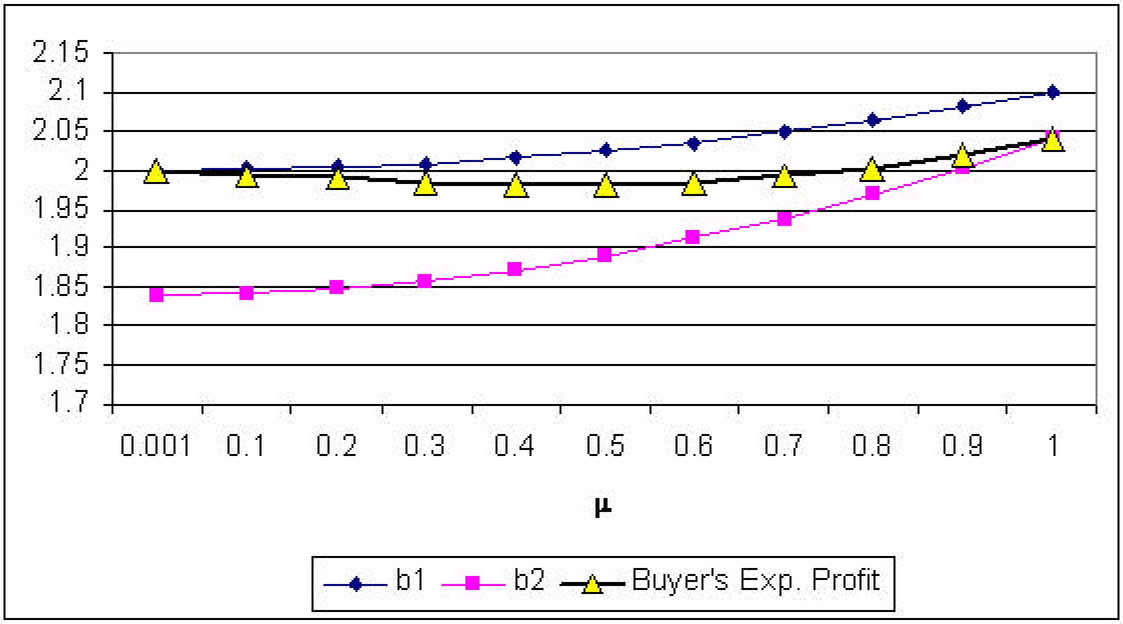

Figure 2: Behavior of buyer's expected profit as a function of μ

8. Comparison with the DSC Model when audit is biased

In this section, we augment the DSC model by including specifically the impact of effort. We then compare the buyer's expected profits under DSC and under audit.

In order to adapt the DSC model to include effort, we modify assumption (5) as follows:

(5a*) Each supplier can exert an unobservable effort level, *e*, that takes values in $\left[\underline{e},\bar{e}\right]$. The cost of effort is given by $E_i(e)$ and its effect is to reduce the unit cost by *e*. Thus, the cost function can be written as

$$C_i(Q,e) = C(Q,q_i,e) = K(q_i) + \int_0^Q c(q,q_i)dq - eQ + E_i(e).$$

(5b*) We further assume that the function $E_i(e)$ is convex and increasing in *e*.

Under those assumptions, we may write the optimal effort as $e^* = g(Q,\ q_i)$. Notice that the buyer is assumed to know this function as a parameter of θ. Thus, we can write the cost as

$$C_i(Q) = C(Q,q_i) = K(q) + \int_0^Q c(q,q_i)dq - g(Q,q_i) + E_i(g(Q,q_i)).$$

Re-define the marginal cost to be,

$$\tilde{c}(q,q_i) = c(q,q_i) - g_Q(Q,q_i) + \partial E_i(g(Q,q_i))/\partial Q \text{ , and}$$

$$C_i(Q) = C(Q,q_i) = K(q_i) + \int_0^Q \tilde{c}(q,q_i)dq.$$

Assume further:

(5c*) The redefined marginal cost is greater than zero for the range of values for Q that we are interested in.

(5d*) The marginal cost is assumed to be increasing and convex in the cost parameter, θ, and $\tilde{c}_{Qq}(Q,q) \geq 0$.

With the above assumptions, and the redefined costs, DSC's assertions are valid. Thus, the optimal buyer's profit under AWA can be computed.

Consider now the option of auction and auditing the firm, i.e., the A&A mechanism. We can directly apply Proposition 5 to determine the optimal value of μ by numerical search. A specific example is presented in the next section.

8.1 A numerical example when audit is biased

We modify the example of section 6.1. Let $E(e, \theta_i) = Ae^2$, where $e \in [0,\bar{e}]$ and A is a positive number. The modified cost function is

$$C(Q,q) = qQ - \frac{Q^2}{4A},$$

where we require A to be greater than or equal to 6 to ensure that the total cost is non-negative and increasing in Q for the entire range of $\theta \in [0.1, 1.1]$ and $Q \in [0,1.1]$. Thus, AWA sets, see (3.2):

$$V'(Q(q) = q - \frac{Q}{2A} + H(q), q \in [\underline{q}, \bar{q}]$$

or

$$2-2Q=q-\frac{Q}{2A}+q-0.1 \Leftrightarrow Q(q)=\frac{1.05-q}{1-\frac{1}{4A}}.$$

The upper cut-off value is $\theta^{**}=1.05$. The expected profit using the AWA mechanism is now given by

$$n\int_{0.1}^{1.05}(1.1-q)^{n-1}\left(2\left(\frac{1.05-q}{1-\frac{1}{4A}}\right)-\left(\frac{1.05-q}{1-\frac{1}{4A}}\right)^2-q\left(\frac{1.05-q}{1-\frac{1}{4A}}\right)-\left(\frac{1.05-q}{1-\frac{1}{4A}}\right)^2/4A-\left(\frac{1.05-q}{1-\frac{1}{4A}}\right)(q-0.1)\right)dq$$

(8.1)

$$=n\int_{0.1}^{1.05}\frac{(1.1-q)^{n-1}(1.05-q)^2}{\left(1-\frac{1}{4A}\right)}dq.$$

Note that the profit achieved is $1/\left(1-\frac{1}{4A}\right)$ times the profits when effort is 0.

Now consider the A&A mechanism. We first solve for the supplier's bid:

(8.2) $$b_i=\max_{Q,e}\left(2Q-Q^2-q_iQ+eQ-\frac{1}{m}Ae^2\right).$$

Solving next for the effort, we get

(8.3) $$e^*=\min\left\{\frac{Qm}{2A},\bar{e}\right\}.$$

Substituting this back into (8.2) and writing the bid as a function of Q:

(8.4) $$b_i(Q)=\left(2Q-Q^2-q_iQ+\frac{Q^2m}{4A}\right)\text{ or}$$

(8.5) $$b_i(Q)=\left(2Q-Q^2-q_iQ+\bar{e}Q-A\bar{e}^2/m\right).$$

However, notice that (8.4) is concave in Q only if μ is less than 4A. Otherwise, it is best to have unbounded Q. But effort is bounded at $\bar{e}$, which prevents Q from being unbounded. Thus,

$$Q^*=\frac{(1-q/2)}{1-\frac{m}{4A}},\ if\ m\le 4A\,and\ \frac{(1-q/2)}{1-\frac{m}{4A}}\frac{m}{2A}\le\bar{e}$$

$$=\left(1-q/2+\bar{e}/2\right)\ else.$$

Substituting these values into (8.2), we obtain the optimal bid and effort pair:

$$(b_i, e^*) = \left(\frac{(1-q/2)^2}{(1-\frac{m}{4A})}, \frac{(1-q/2)}{1-\frac{m}{4A}}\frac{m}{2A}\right) \text{ if } m \le 4A \text{ and } \frac{(1-q/2)}{1-\frac{m}{4A}}\frac{m}{2A} \le \bar{e}$$

$$= ((1-q/2+\bar{e}/2)^2 - A\bar{e}^2/m, \bar{e}) \text{ otherwise.}$$

These expressions can now be used to determine the expected profit for the buyer, and thus the optimal choice of μ. The comparison of AWA versus A&A is shown in Table 2. The table utilizes the value 0.01 to represent the smallest possible level of μ.

	AWA	A&A							
n	$E[\Pi^{B*}]$	μ^*	d_1	d_2	$E_1(e)$	Buyer's Exp. Profit	% inc.	Channel Exp. Profit	First Best Exp Profit
2	0.4470	0.0100	0.6273	0.3940	0.0000	0.6252	39.86%	0.6275	0.6436
5	0.6497	0.0100	0.7582	0.6229	0.0000	0.7571	16.53%	0.7585	0.7755
10	0.7640	0.0100	0.8202	0.7419	0.0000	0.8196	7.27%	0.8204	0.8409
20	0.8373	0.0595	0.8603	0.8169	0.0001	0.8590	2.59%	0.8616	0.8804
30	0.8648	0.4060	0.8813	0.8514	0.0037	0.8746	1.13%	0.8867	0.8947
40	0.8793	0.5545	0.8923	0.8696	0.0070	0.8853	0.68%	0.8979	0.9022
50	0.8882	0.6040	0.8977	0.8792	0.0083	0.8920	0.43%	0.9032	0.9067
100	0.9065	0.8020	0.9114	0.9020	0.0150	0.9076	0.12%	0.9151	0.9160

Table 2: Buyer's Expected Profit using AWA and A&A (A = 10, $\bar{e}$ = 0.2)

Observing the table we see that for small values of *n*, $\mu^* = 0$, i.e., we rely heavily on the audit. This is because the difference between the first and second bidder tends to be large. Thus, for the buyer, the cost of inducing the supplier to exert optimum effort is too large. The buyer opts to expropriate the entire channel profits, and the supplier in turn exerts minimal effort. In the case of large *n*, the difference between the first and second bidder tend to be miniscule, and thus the buyer is ready to give up on this difference in order to ensure maximal effort. At large values of *n*, the increase in the buyer's expected profit due to A&A over AWA is also small. This further emphasizes the reliance on the auction rather than the audit for large *n*.

	Auction above n = 40, audit otherwise						
n	μ^*	d_1	d_2	$E_1(e)$	Buyer's Exp. Profit	% Inc. over AWA	% loss versus A&A
2	0.01	0.6273	0.3940	0.0000	0.6252	39.86%	0.00%
5	0.01	0.7582	0.6229	0.0000	0.7571	16.53%	0.00%
10	0.01	0.8202	0.7419	0.0000	0.8196	7.27%	0.00%
20	0.01	0.8584	0.8158	0	0.8582	2.50%	-0.09%
30	0.01	0.8727	0.8431	0	0.8727	0.91%	-0.22%
40	1	0.9025	0.8796	0.0231	0.8796	0.03%	-0.64%
50	1	0.9067	0.8878	0.0232	0.8878	-0.05%	-0.47%
100	1	0.9162	0.9065	0.0235	0.9065	-0.01%	-0.12%

Table 3: Buyer’s Expected Profit using naïve choice

Finally, in keeping with the theme of the paper, we must ask what happens if the buyer does not have sufficient information to determine the optimal value of μ. Motivated by the discussion above, the buyer could use a naïve choice: use audit as the primary mechanism for *n* less than or equal to 30, and use auction for other cases. The resulting expected profit to the buyer are shown in Table 3. As expected the naïve choice does not reduce the expected profit greatly.

9. Discussion

We have examined supply chain auctions from the perspective of information available to a buyer and her suppliers. We have presented an extremely general model of 2-stage supply chains and have examined mechanisms available to the buyer to optimize her profits, given her partial information about suppliers’ costs. We have shown that under a minimal set of assumptions, the buyer can induce efficient outcomes, and at the same time expropriate reasonable level of profit for herself. Moreover, we have shown that audits provide a practical mode of extracting information about supply chain partners, which allows the buyer to expropriate any desired fraction of the total (centralized) channel profits. Finally, we have examined the behavior of audit-based auctions when the audit is biased. When the number of suppliers is small, audit with auction seems to be preferred to isolate the best supplier. When the number of suppliers is large, auction is preferred.

A number of extensions and generalizations to the problem studied in this paper suggest themselves. These include studying the core of the cooperative game alluded to in Proposition 3, extending the analysis to several buyers and suppliers, and to one buyer with contract award to several suppliers. Empirical research questions based on this paper include: When is one or the other form of mechanism preferred by buyers and suppliers? What data is collected by buyers about their suppliers, viz., what is the nature of supply intelligence gathered about potential suppliers? What incentive contracts are offered for single versus repeat purchasing situations? How are profit sharing contracts structured and what is the role of bargaining power as well as industry norms in determining the structure of such contracts?

References

Anupindi, R. and Y. Bassok, " Supply Contracts", in Tayur S., R. Ganeshan and M. Magazine (eds.), *Quantitative Models for Supply Chain Management*, Kluwer, Boston, MA, 1998

Baron, D. and D. Basanko. 1984. Regulation, asymmetric information and auditing. *Rand J. of Economics* 15:447-470

Cachon, G., "Chapter 6: Supply Chain Coordination with Contracts," in forthcoming, *Handbook in Operations Research and Management Science: Supply Chain Management*, edited by Steve Graves and Ton de Kok, North-Holland Publishers, 2003.

Chen, F., "Auctioning Supply Contracts," working paper Graduate School of Business, Columbia University, New York, NY 10027, 2001 (revised March 2003).

Chen F, "Information Sharing and Supply Chain Coordination, To appear in the Handbooks in Operations Research and Management Science: Supply Chain Management, edited by Ton G. de Kok and Stephen C. Graves, 2002.

Dasgupta, S. and Spulber, D. F., “Managing Procurement Auctions,” *Information Economics and Policy*, **4**, 1990, 5-29.

Elmaghraby, Wedad J., “Supply Contract Competition and Sourcing Policies,” *Manufacturing & Service Operations Management*, **2**, 4, 2000, 350-371.

Hansen, R., “Auctions with Contingent Payment,” *American Economic Review*, **75**, 1985, 82-65.

Hayek, F. A. 1945. The use of knowledge in society. *American Economic Review*. 4: 519-530.

Klemperer, P., “Auction Theory: A guide to the Literature,” *Journal of Economic Surveys*, **13**, 3, 1999, 227-86.

Laffont, J. and J. Tirole, *A Theory of Incentives in Procurement and Regulation*, The MIT Press, Cambridge, MA, 1994.

Maskin, E. and J. Riley, “Optimal Auctions and Risk Averse Buyers,” *Econometrica*, **52**, 1984, 1473-1518.

McMillan, J. and R. P. McAfee. 1987. Auctions and Bidding. *Journal of Economic Literature*. 25:699-738.

Milgrom, P., “Rational Expectations, Information Acquisition, and Competitive Bidding,” *Econometrica*, **49**, 4, 921-943.

Milgrom, P., “Auctions and Bidding: A Primer,” *Journal of Economic Perspectives*, **3**, 3, 1989, 3-22.

Myerson, Roger, “Optimal Auction Design,” *Mathematics of Operation Research*, **6**, 1981, 58-73.

Pasternack, B., “Optimal pricing and returns policies for perishable commodities,” *Marketing Science*, **4**, 2, 1985, 166-76.

Riley, J. and W. Samuelson, “Optimal Auctions,” *American EconomicReview*, **71**, 381-392.

Samuelson, W., “Auctions with Contingent Payments: Comment,” *American Economic Review*, **77**, 4, 1987, 740-745.

Seshadri, S., “Bidding for Contests,” Management Science, 41, 4, 1995, 561-576.

Seshadri, S. and E. Zemel, “The core of supply chain games,” under preparation, Stern School of Business, New York University, 2003.

Tayur, S., R. Ganeshan and M. Magazine (eds.), *Quantitative Models for Supply Chain Management*, Kluwer, Boston, MA, 1998

Townsend, R. M., “Optimal Contracts and Competitive Markets with Costly State Verification,” *Journal of Economic Theory*, **21**, 1979, 265-293.