

Combinatorial Auctions for Supply Chain Formation

William E. Walsh Michael P. Wellman
University of Michigan
Artificial Intelligence Laboratory
1101 Beal Av, Ann Arbor, MI 48109-2110
USA
wew, wellman @umich.edu

Fredrik Ygge
EnerSearch AB and Uppsala University
Chalmers Science Park
SE-412 88 Gothenburg
Sweden
ygge@enersearch.se

ABSTRACT

Supply chain formation presents difficult coordination issues for distributed negotiation protocols. Agents must simultaneously negotiate production relationships at multiple levels, with important interdependencies among inputs and outputs at each level. Combinatorial auctions address this problem by global optimization over expressed offers to engage in compound exchanges. A one-shot combinatorial auction that optimizes the reported value of the bids results in optimal allocations with truthful bids. But autonomous self-interested agents have an incentive to bid strategically in an attempt to gain extra surplus. We investigate a particular combinatorial protocol consisting of a one-shot auction and a strategic bidding policy. We experimentally analyze the efficiency and producer surplus obtained in five networks, and compare this performance to that of a distributed, progressive auction protocol with non-strategic bidding. We find that producers can sometimes gain significantly by bidding strategically. However, when the available surplus is small relative to the consumers' values, the producers' strategic behavior may prevent the supply chain from forming at all, resulting in zero gains for all agents. We examine the robustness of the combinatorial protocol by investigating agent incentives to deviate, identifying quasi-equilibrium behavior for an example network.

1. INTRODUCTION

Complex business negotiations often involve interrelated exchange relationships among multiple levels of production, often referred to as a *supply chain*. To respond to rapidly changing market conditions, companies must be able to dynamically form and dissolve business interactions, requiring automated support for *supply chain formation*, the process of bottom-up assembly of complex production and exchange relationships.

The problem of supply chain formation can be particularly challenging when firms must contend for scarce resources at multiple levels in the supply chain hierarchy [15]. That resource availability cannot be guaranteed until contracts are finalized makes it difficult for firms to coordinate the negotiations between their various

input resources and their production outputs. Production technologies typically contain strong *complementarities*. That is, the values of obtaining various inputs and for producing outputs are mutually dependent. A firm could be penalized (either explicitly, or by loss of reputation) if it is *infeasible*, that is unable to acquire all inputs necessary to meet its production obligations, and could be unprofitable if it acquires extraneous inputs or produces goods it cannot sell.

In prior work [13, 14], we studied a distributed, progressive market protocol with non-strategic bidding—described in detail below—for supply chain formation. Despite the fact that agents negotiate for goods separately using simple, localized bidding policies, they can often effectively form high-quality supply chains by following the protocol. However, there exist situations in which the complementarities cause the protocol to fail to form the optimal supply chain. Although the protocol guarantees that agents are never infeasible and are always profitable when they sell their output, producers can lose significant profits by acquiring inputs but failing to sell outputs.

To address the problem of complementarities, researchers have proposed mechanisms that mediate the negotiation of several interdependent goods through a single entity [7, 10, 17]. In *combinatorial auctions*, agents place all-or-nothing bids for bundles of goods. The auction computes a high-quality allocation of bundles, ensuring that agents do not receive undesirable partial bundles. Although the general problem of optimal winner determination is NP-hard [11], Anderson et al. [1] have shown that a commercial mixed-integer-linear programming package can quickly solve many large allocation problems. In the present work we study a combinatorial auction protocol, consisting of a one-shot combinatorial auction and strategic bidding policy, in production economies for supply chain formation.

In the next section, we describe our model of supply chain structures as task dependency networks. In Section 3 we review our previous work in market-based supply chain formation. We present our combinatorial protocol in Section 4. We discuss some fundamental issues in comparing protocols in Section 5, and describe experiments comparing the combinatorial protocol with the distributed protocol in Section 6. In Section 7 we examine agents' incentives to deviate from the combinatorial protocol and identify approximate, constrained strategic equilibria for an example problem. We conclude and discuss future work in Section 8.

2. TASK DEPENDENCY NETWORKS

Supply chain formation is, informally, the problem of assembling a network of agents that, given local knowledge and communication, can transform basic goods into composite goods of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'00, October 17-20, 2000, Minneapolis, Minnesota.
Copyright 2000 ACM 1-58113-272-7/00/0010 ..\$5.00

value [15]. In our model, we use the term “good” to refer to any discrete resource or task. We further assume that goods are rival, meaning that individual units cannot be shared between agents.

We choose an abstract model which glosses over some aspects of supply chain formation but emphasizes the understudied problem of resource contention at multiple levels of the supply chain hierarchy. More precisely, we formulate the problem as follows (a variant of a model from our previous work [14]). A **task dependency network** is a directed, acyclic graph, (V, E) , representing dependencies among agents and goods. $V = G \cup A$, where G is the set of goods and $A = C \cup \Pi$ is the set of agents, comprised of consumers C , and producers Π . Edges, E , connect agents with goods they can use or provide. There exists an edge $\langle g, a \rangle$ from $g \in G$ to $a \in A$ when agent a can make use of one unit of g , and an edge $\langle a, g \rangle$ when a can provide one unit of g . When an agent can use or provide multiple units of a good, separately indexed edges represent each unit. The goods can be traded only in discrete quantities.

A **consumer** c wishes to acquire one unit of a particular good and obtains **value** v_c for doing so. A **producer** π can produce a single unit of an **output** good conditional on acquiring a certain number each of some fixed set of zero or more **input** goods. π must acquire each of its inputs, and in addition incurs **cost** κ_π to provide its output.¹ The cost might represent the value π could obtain from alternative, mutually exclusive activities not explicitly modeled, or some actual direct cost of its production activity.

We express values and costs in monetary units. Figure 1 shows an example network with specific consumer values.




Figure 1: Task Dependency Network NO-EQUILIB (name to be explained). Circles represent goods, rectangles and octagons represent agents, and arrows represent use or provision of goods. Dollar amounts indicate consumer values.

An **allocation** is a subgraph $(V', E') \subseteq (V, E)$. For $g \in G$, an edge $\langle a, g \rangle \in E'$ means that agent a provides g , and $\langle g, a \rangle \in E'$ means a acquires g . An agent is in an allocation graph iff it acquires or provides a good. A good is in an allocation graph iff it is acquired or provided.

A producer is **active** iff it provides its output. A **producer is feasible** iff it is inactive or acquires all its inputs. Consumers are always feasible. An **allocation is feasible** iff all agents are feasible and all goods are in **material balance**, that is the number of edges into a good equals the number of edges out. Consumer c obtains value $v_c(E')$ for allocation E' , which is v_c if it obtains its desired good, otherwise zero. Producer π incurs cost $\kappa_\pi(E')$ which is κ_π if it provides its output, otherwise zero.

¹Those familiar with our previous work [13, 14] will observe that we are effectively merging “suppliers” and “producers” into a single, more general producer type of agent. Here we allow producers to have fixed costs, and model a “supplier” as a producer with no inputs.

DEFINITION 1 (VALUE OF AN ALLOCATION). The value of allocation (V', E') is:

$$\text{value}((V', E')) \equiv \sum_{c \in C} v_c(E') - \sum_{\pi \in \Pi} \kappa_\pi(E').$$

DEFINITION 2 (EFFICIENT ALLOCATIONS). The set of efficient allocations contains all feasible allocations (V^*, E^*) such that $\text{value}((V^*, E^*)) = \max_{(V', E') \text{ feasible}} \{\text{value}((V', E')) \mid (V', E') \text{ is feasible}\}$.

A **solution** is a feasible allocation such that one or more consumers acquires a desired good. If $c \in C \cap V'$ for solution (V', E') , then (V', E') is a **solution for consumer** c . Figure 2 shows a solution and efficient allocation for the network of Figure 1 (with some specific costs).




Figure 2: An efficient solution for Network NO-EQUILIB, given costs as shown. Shaded goods and agents and solid edges are in the solution.

3. A SEPARATELY-PRICED-GOODS APPROACH

In previous work, we examined a market approach to supply chain formation that priced goods separately. In this approach, we posit a **price system** p , which assigns to each good g a nonnegative number $p(g)$ as its **price**. Intuitively, prices indicate the relative value of the goods. Therefore, agents may use the prices as a guide to their local decision making.

Informally, a (**competitive**) **price equilibrium** is a feasible allocation in which every agent is optimizing with respect to prices. Agents optimize **competitively** (non-strategically) in that they take prices as given, ignoring any effects of their own actions on these prices. In a price equilibrium allocation for a task dependency network, participating consumers acquire goods for prices no greater than their values, producers make nonnegative profits, and consumers and producers not in the allocation would not obtain positive surplus at the corresponding prices. Price equilibria are desirable because they are stable by definition and efficient in well-behaved economies. Given quasilinear utility, price equilibria have been shown to be efficient under fairly general conditions [2, 3, 19], and in task dependency networks in particular [14]. However, price equilibria may not exist given discrete goods and complementarities. For example, there is no price system supporting a price equilibrium for the network of Figure 2, given the costs shown (hence the name NO-EQUILIB). Moreover, we do not know of any market protocol based on prices for separate goods that guarantees convergence to price equilibria under general conditions, even when they exist, or that guarantees any bounds on the quality of the allocation.

Nevertheless, we have identified a separately-priced-goods market protocol, Simultaneous Ascending (M+1)st Price with Simple

Bidding (SAMP-SB), that often performs well [13, 14]. In the protocol, agents negotiate simultaneously in separate ascending auctions, one for each good. We analyzed the protocol assuming that agents use non-strategic myopic bidding policies, relying only on the price reports from the auctions for their own goods of interest. Under these assumptions, SAMP-SB is guaranteed to produce feasible allocations [13].

SAMP-SB is not guaranteed to produce efficient allocations. We categorize suboptimality into three different (not necessarily exclusive) types:

1. Failure to find a solution when an efficient solution exists. This can occur when there is insufficient slack between consumer values and producer costs.
2. Forming a solution with value inferior to an efficient solution.
3. Inactive producers acquiring inputs.

Experiments on random networks suggest that Type 3 is the dominant type of suboptimality [14], although Types 1 and 2 can and do occur.

Figure 3 shows the results of a run of SAMP-SB on the network of Figure 2, producing a suboptimal allocation. The allocation exhibits Type 2 suboptimality because the efficient allocation would include consumer $c2$ but not $c1$. Note also that the costs of producers $a2$ and $a4$ do not contribute to obtaining any value in the network, an instance of Type 3 of suboptimality.




Figure 3: A suboptimal allocation generated by SAMP-SB in a network without a price equilibrium. Shaded agents and goods and solid lines are in the allocation.

One way to eliminate Type 3 suboptimality is to allow *decommitment*. SAMP-SB-D is the protocol whereby agents follow SAMP-SB, but inactive producers subsequently *decommit* from their inputs. That is, they are released from all exchange obligations determined by the auctions. The decommitment proceeds recursively until all inactive producers decommit from their inputs. For example, in Figure 3, agents $a4$ and $a2$ would decommit. Decommitment has the benefit that agents are ensured to be no worse off by participating in SAMP-SB-D. However, this is achieved by making the auction allocations non-binding, which is undesirable to the producers who lose their output sales to decommitments. It also begs the question of how to enforce the requirement that inactive producers be the only agents that decommit.

Although suboptimal and even undesirable (to some producers) allocations can result, many market environments support only separate, distributed negotiations. Ultimately, supply chain formation is economically intertwined with a broader market that is too large to feasibly coordinate with a single, unified negotiation. That noted, in many industries we can identify smaller, strongly intradependent submarkets that are relatively weakly dependent on the broader

market. In such environments it is possible for a “market maker” or a consortium to support a single mechanism that coordinates all of the most important activities within a market. Agents can then engage in all-or-nothing negotiations for their goods of interest to ensure that they are not stuck with undesirable partial bundles of goods. In the next section we describe a combinatorial auction protocol that, by directly linking the negotiations for all goods, ameliorates some of the coordination problems that occur in protocols built on separate negotiations.

4. COMBINATORIAL PROTOCOL

To address the shortcomings of SAMP-SB, and of separately-priced-goods approaches in general, we describe a particular combinatorial protocol, consisting of an auction mechanism and agent bidding policies. Agents submit all-or-nothing bids for bundles of goods to a single combinatorial auction. This prevents Type 3 suboptimality by ensuring that inactive producers do not acquire inputs, as can occur in SAMP-SB.

An *auction* receives all bids, enforces bid rules, and reports information to agents and computes allocations as a function of bids received. Agents utilize *bidding policies* that specify how to bid in the auctions as a function of their preferences and history of messages received from the auctions. The key distinction between the auction mechanism and bidding policies is that the former is under the control of the system designers, whereas the latter are determined by individual agents. The protocol as a whole is the subject of our analysis.

4.1 Combinatorial Auction Mechanism

The combinatorial auction we study is a one-shot mechanism: agents submit bids reporting costs and values, then the auction computes an allocation that maximizes the reported value and informs the agents of results. An agent pays the price it bid for the allocation it receives. If the auction receives more money than it pays out, the proceeds are distributed evenly among *all* consumers.² Agent a places a bid b_a of the form

$$\langle r_a, \langle g_1, q_a^1 \rangle, \dots, \langle g_n, q_a^n \rangle \rangle,$$

where q_a^i is the integer quantity that agent a demands (positive for input demands and negative for output demands) for good g_i , and r_a is its reported willingness to pay (or be paid, in the case of negative numbers) for the demanded bundle of goods. For instance, a producer that requires one unit each of inputs g_1 and g_2 to produce output g_3 and seeks a payment of 5, would place the bid $\langle -5, \langle g_1, 1 \rangle, \langle g_2, 1 \rangle, \langle g_3, -1 \rangle \rangle$.

Given a set of bids B , the auction computes the winning allocation from:

$$\Psi(B) = \max_{\mathbf{x}} \sum_{b_a \in B} r_a x_a \text{ s.t. } \sum_{b_a \in B} q_a^i x_a = 0, i = 1 \dots n,$$

where $x_a = 1$ if agent a wins its bid, and $x_a = 0$ otherwise. This guarantees global feasibility as long as each producer’s bid represents a locally feasible combination of inputs and outputs. Moreover, this computes the allocation that optimizes value *as reported*

²Although it is possible to define “Generalized Vickrey” payments that induce truth revelation as a dominant strategy [5], such a scheme will in general require subsidies in the supply-chain context. The mechanism defined here is guaranteed to be budget balanced. It is an interesting open question whether alternative designs would have positive incentive properties. The fact that agent preferences have a very regular form—reminiscent of the “single-minded” assumption explored by Lehmann et al. [4]—may present opportunities for mechanism designers to exploit.

by the bids. Note that we assume producers have a single possible set of inputs. We would need a somewhat more complex equation for $\Psi(B)$ if producers had alternative production possibilities.

Agent a 's *surplus* is the excess value it receives from the allocation computed by the auction. Producer π gets surplus $x_\pi(r_\pi - \kappa_\pi)$. We divide the remaining reported surplus equally among all consumers. Hence, consumer c gets surplus $x_c(v_c - r_c) + \Psi(B)/|C|$. The total surplus available in a network is equal to the value of an efficient allocation.

4.2 Combinatorial Bidding Policies

Although multiagent system designers do not generally have control over the agents' behaviors, any conclusions about the outcome of a protocol must be based on some assumptions about these behaviors. We can see immediately that if agents behave non-strategically (i.e., bid their true valuations) in the combinatorial auction mechanism, then the result will always be an efficient allocation.³ Hence, in this work we choose to study the performance of strategic bidding policies. We also propose that, because the combinatorial auction eliminates the risk of agents acquiring unprofitable bundles of goods, agents may be more willing to place aggressive, strategic bids than they might when faced with separate negotiations.

To structure our analysis, we assume that it is common knowledge that the consumers bid their true values, and limit our attention to the producers' strategic behavior. The fact that we distribute excess surplus evenly among all consumers does reduce the incentives for consumers to behave strategically. But more to the point, we can assume a non-strategic class of consumer agents without loss of generality, as we could model a strategic consumer as follows. The consumer places a non-strategic bid for a dummy good and also places a bid as a producer. This producer can uniquely provide the dummy good but requires as input the good that the consumer truly wants.

We assume common knowledge of the structure of the network and consumers' values, and that producers' costs are drawn from a probability distribution uniform on $[0, 1]$. A common strategic model in the auction literature assumes that agents play Bayes-Nash equilibrium strategies based on common knowledge of cost and value distributions. Consider the degenerate (no task dependency) network of N parallel producers and one consumer with unit value, shown in Figure 4. We can derive the Bayes-Nash equilibrium for the producers in this network by recognizing its correspondence to the first-price sealed-bid auction. McAfee and McMillan [6] show that in this auction type, if N buyers have valuations for a single common good, chosen from probability distributions uniform on $[0, 1]$, and buyer i has value v_i for the good, then a Bayes-Nash equilibrium bidding policy is $v_i - (\int_0^{v_i} x^{N-1} dx) / v_i^{N-1}$. Adjusting the notation and signs to our context, this becomes, for producer π ,

$$r_\pi = -\kappa_\pi - \frac{\int_{\kappa_\pi}^1 (1-x)^{N-1} dx}{(1-\kappa_\pi)^{N-1}},$$

which reduces to

$$r_\pi = -\kappa_\pi - \frac{1}{N}(1-\kappa_\pi). \quad (1)$$

³The conclusion is not so immediate—or even true—for protocols based on individual goods, such as SAMP-SB. The same holds for iterative combinatorial mechanisms [8, 18], where agents make incremental offers for only some bundles of potential interest. Determining the implications of various strategic and non-strategic bidding policies for these mechanisms remains an interesting and complex problem.

Observe that, according to (1), agents are strategically offering to accept payments above their cost, and that this premium converges to zero as the number of producers grows.




Figure 4: A network with a closed-form Bayes-Nash equilibrium.

Unfortunately, deriving such Bayes-Nash equilibria is difficult for even slightly more complicated networks. Because we cannot produce closed-form Bayes-Nash equilibrium expressions for any of the other networks investigated in this paper, we instead attempt to define a plausible strategic bidding policy. In particular, we assume that a producer bids to obtain a fraction of the expected available surplus scaled by the expected proportion of its contribution to the global value. We show below that this bidding policy is a generalization of the Bayes-Nash equilibrium policy for the simple parallel network in Figure 4, in that the general policy reduces to (1) for that network.

Let $\hat{\Psi}(Y)$ denote the surplus available from agents $Y \subseteq A$, equivalent to $\Psi(B)$ for the case where B consists of bids corresponding to true values and costs of agents Y . Let $\Pi^* \subseteq A$ be the producers participating in the efficient allocation. The contribution of these producers to the value of the allocation is Δ^* , where

$$\Delta^* = \hat{\Psi}(A) - \hat{\Psi}(A \setminus \Pi^*).$$

The contribution Δ_π of a producer π to the value of an allocation is the difference between the efficient global value, and the global value with π excluded from the allocation,

$$\Delta_\pi = \hat{\Psi}(A) - \hat{\Psi}(A \setminus \{\pi\}).$$

Note that $\Delta_\pi = 0$ for π not part of an efficient allocation.

A producer's relative contribution can then be defined in terms of its expected proportional contribution, conditional on its being part of the efficient allocation. We assume that producer π bids

$$r_\pi = -\kappa_\pi - E \left[\frac{\Delta_\pi \Delta^*}{\sum_{\hat{\pi} \in \Pi} \Delta_{\hat{\pi}}} \middle| \pi \in \Pi^* \right]. \quad (2)$$

If π cannot possibly participate in the efficient allocation, the expectation term of (2) is undefined, and it does not bid.

For the pure parallel network of Figure 4, it turns out that this bidding policy corresponds to the Bayes-Nash equilibrium strategy. To see this, first observe that for this network, Π^* consists of exactly one producer. When this producer is π , $\Delta_\pi = \sum_{\hat{\pi} \in \Pi} \Delta_{\hat{\pi}}$, so we need only compute

$$r_\pi = -\kappa_\pi - E[\Delta^* | \pi \in \Pi^*].$$

Letting $X_{(i)}$ denote the i th order statistic (i th lowest) of producer costs, this reduces to

$$r_\pi = -E[X_{(2)} | X_{(1)} = \kappa_\pi]. \quad (3)$$

Equation (3) is equivalent to the first order statistic of the remaining $N - 1$ producer costs (excluding κ_π) given $X_{(1)} = \kappa_\pi$. But also given $X_{(1)} = \kappa_\pi$, each of the remaining producer costs is

uniformly distributed on $[\kappa_\pi, 1]$. Therefore, we can compute (3) using (1).⁴

Though it is reassuring that we can justify this bidding policy for the case we can solve, this clearly does not entail any firm conclusions for more complicated networks. The bidding policy has a note of plausibility in that winning producers proportionally share in their expected contribution to the global value. However, we can empirically show that the bidding policy is not a Bayes-Nash equilibrium. In Section 7, we discuss how a producer might deviate from this base bidding policy and explore the possibility that variants of the policy may constitute a “quasi-equilibrium”.

5. COMPARING PROTOCOLS

Figure 5 illustrates one way to divide the space of protocols into four (not necessarily exhaustive) classes. The two columns distinguish mechanisms by scope: whether all issues are resolved by a global combinatorial auction, or by a distributed collection of individual-good auctions. The two rows distinguish agent bidding policies according to whether or not they are strategic, that is, take into account the agent’s own effect on allocation and monetary transfers.

	Combinatorial	Distributed
Non-strategic	one-shot combinatorial/ true costs	SAMP-SB
Strategic	one-shot combinatorial/ Equation (2)	?

Figure 5: A classification of the protocols we consider. The columns distinguish mechanism scope and the rows distinguish whether bidding policies are strategic.

As noted above, non-strategic bidding in the combinatorial mechanism produces perfectly efficient allocations, virtually by definition. This is a natural benchmark by which to judge performance of other protocols. Our experiments evaluate the *efficiency*—fraction of the efficient value obtained, of the strategic combinatorial protocol we develop in Section 4.

Our representative of the class of distributed protocols is SAMP-SB, selected primarily on the basis of our previous studies. Because it differs in both scope (distributed versus combinatorial) and strategy, the overall qualitative relationship between the two protocols is ambiguous. We would expect that the wider scope of the combinatorial mechanism would promote efficiency, whereas strategic behavior might impede it. This renders the comparison an interesting subject of experimental analysis.

We must emphasize that the object of such a performance comparison cannot be to determine which protocol is “better”—if this is even a meaningful question. The mechanism designer must consider factors beyond allocative efficiency, such as computational costs, communication latencies, and authority over negotiation scope. Perhaps more critically, agent bidding behaviors are beyond the mechanism designer’s control, and so imposing either of these protocols is literally impossible.

Still, to make progress in understanding a challenging problem, we would like to compare the most plausible protocols associated with particular mechanism choices. We have argued the plausibility of a particular bidding policy in the combinatorial mechanism

⁴This is not surprising, given that the Revenue Equivalence Theorem [6] implies an analogous result for the first-price sealed-bid auction.

(while acknowledging that it does not generally constitute a Bayes-Nash equilibrium), and put forth some (probably less compelling) reasons that one might pursue simple, non-strategic, myopic policies in the face of a progressive distributed mechanism.⁵ Nevertheless, perhaps the best explanation for the emptiness of the bottom-right cell in Figure 5 is that coming up with a justifiable strategic behavior for the distributed progressive case is quite difficult. This is attributable at least as much to dynamic complexity as it is to distributivity, and thus would apply as well to iterative combinatorial auctions [8, 9, 18].

6. BASE PROTOCOL EXPERIMENTS

In order to obtain an understanding of the effectiveness of the combinatorial auction protocol we propose, we simulated the protocol on five task dependency networks, and compared the quality of the allocations and producer surplus to SAMP-SB. We describe the experimental setup in Section 6.1, the implementation of the experiments in Section 6.2, and the results in Section 6.3.

6.1 Setup

For this investigation, we chose to focus on a small set of networks exhibiting a variety of structural properties: NO-EQUILIB (Figure 1), SIMPLE (Figure 6), UNBALANCED (Figure 7), BIG (Figure 8), and MANY-CONS (Figure 9). For each consumer in a network, we calculated a fixed value so that, excluding all other consumers, there exists a positive-surplus solution for this consumer with 0.9 probability. We determined consumer values via simulation, assuming the specified distributions of producer costs.




Figure 6: Network SIMPLE.




Figure 7: Network UNBALANCED.

We ran experiments on multiple instances of each network. For each instance we randomly chose producer costs uniformly over $[0, 1]$, and for each producer, based its bid on 500 Monte Carlo samples over the costs of other producers in its instance. We also

⁵It is well established that pursuing aggressively strategic behavior can lead to pitfalls in uncertain environments [12, 16]. Proposers of particular bidding behaviors (strategic or not) in particular environments still face the burden of demonstrating reasonableness.




Figure 8: Network BIG.

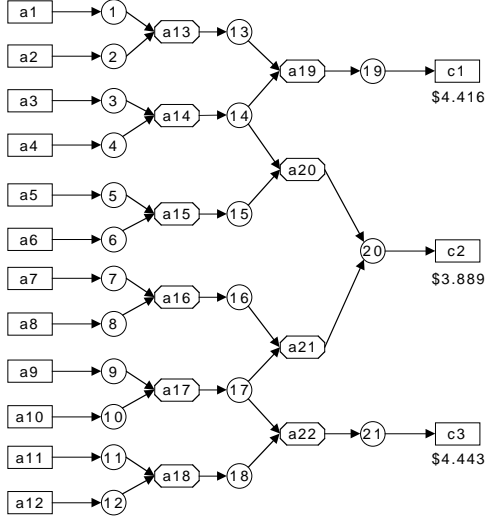


Figure 9: Network MANY-CONS.

ran SAMP-SB and SAMP-SB-D for each instance. For each instance and each protocol, we measured the efficiency as well as the fraction of the optimal value obtained as surplus by the producers.

6.2 Implementation

Lacking a general closed form for the agents' bidding policy, in our experiments we employ Monte Carlo simulation to compute approximations to (2). In each Monte Carlo sample, producer π generates costs for the other producers uniformly from $[0, 1]$. If π is in an efficient allocation given the costs, it computes Δ^* and Δ_π , based on its true cost, the consumers' true values, and the costs it randomly selected for the other producers. From these, it computes \tilde{r}_π , its "sample reported willingness to pay":

$$\tilde{r}_\pi = -\kappa_\pi - \frac{\Delta_\pi \Delta^*}{\sum_{\hat{\pi} \in \Pi} \Delta_{\hat{\pi}}}$$

After running the simulations, π computes its final r_π for its bid as the average of all computed \tilde{r}_π values. If it does not participate in the efficient allocation for any sample, it refrains from bidding.

To solve the combinatorial auction allocation (winner determination) problem of $\Psi(B)$, we used a commercial mixed-integer-linear programming package (CPLEX), an approach introduced by Andersson et al. [1]. We also used the same approach to calculate the Δ_π and Δ^* values for the simulations (which are essentially equivalent to winner determination, but with different parameters). CPLEX solved each optimization in a fraction of a second, and the

total runtime for simulations seemed to be dominated by communications between CPLEX and the Perl script wrapper we employed to manage and aggregate the data.

6.3 Results

We found that the efficiency of the combinatorial protocol was strongly dependent on the available surplus, $\Psi(A)$. Figures 10–15 present the efficiency achieved by the combinatorial protocol on our series of example networks. Each instance is plotted as a point with available surplus measured on the horizontal axis, and the efficiency of the combinatorial auction allocation shown on the vertical axis. Above the plots are shaded regions showing abstractions of the results into classes of efficiency as a function of available surplus.

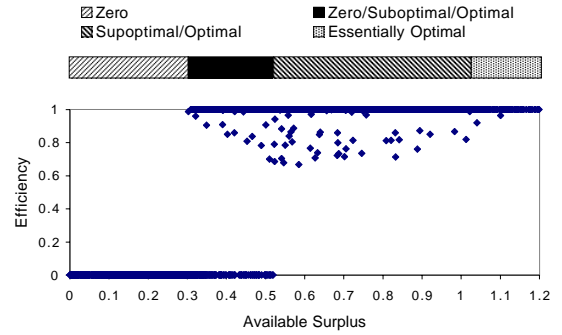


Figure 10: Results of combinatorial protocol experiments for Network SIMPLE. Number of instances: 3220.




Figure 11: Results of combinatorial protocol experiments for Network UNBALANCED. Number of instances: 100.

For all networks, the combinatorial protocol efficiency was always zero for some region in the low range of the available surplus. For all networks—except NO-EQUILIB conditional on nonequilibrium (Figure 13)—there were "Essentially Optimal" regions (most instances optimal, with a few slightly suboptimal instances) at the high range of available surplus. Most networks also exhibited regions of mixed efficiency classes in the middle regions of available surplus.

In contrast to the combinatorial protocol, we did not observe a strong relationship between the efficiency classes of SAMP-SB or SAMP-SB-D and the available surplus. Nevertheless, we did observe distinct and varied efficiency classes. Table 1 shows the distribution of efficiency classes for each protocol and each network.

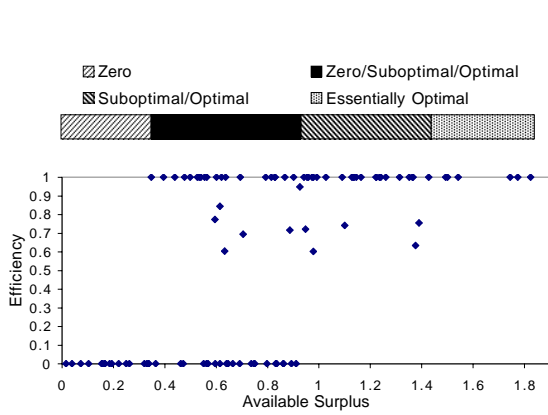


Figure 12: Results of combinatorial protocol experiments for Network NO-EQUILIB conditional on cost structures for which price equilibrium exists. Number of instances: 100.




Figure 13: Results of combinatorial protocol experiments for Network NO-EQUILIB conditional on cost structures for which no price equilibrium exists. Number of instances: 100.




Figure 14: Results of combinatorial protocol experiments for Network BIG. Number of instances: 100.




Figure 15: Results of combinatorial protocol experiments for Network MANY-CONS. Number of instances: 100.

We can infer the number of instances exhibiting Type 3 suboptimality in SAMP-SB by examining the differences between SAMP-SB-D and SAMP-SB totaled over the Negative, Zero, and Suboptimal columns. Decommithment does not affect the contribution of Types 1 and 2 suboptimality, but helps reveal them by eliminating Type 3 suboptimality. Hence, we can infer the number of instances exhibiting Types 1 and 2 suboptimality in SAMP-SB by examining the Zero and Suboptimal columns of SAMP-SB-D, respectively.

Tables 2 and 3 show the average efficiency and average fraction of available surplus obtained by producers, respectively, in each network, for the three protocols. Recall that extra surplus (available as specified by the strategic bids) not taken by producers is distributed evenly among all consumers.

In summary, we found that producers can obtain significant positive surplus with the strategic bidding policy. However, when the available surplus is small relative to the consumers' values, the protocol generates low global efficiency compared to SAMP-SB-D, and in some cases, even SAMP-SB.

The combinatorial protocol solves the problem of Type 3 suboptimality by design, which can produce significant benefits in some networks, as compared to SAMP-SB. Although decommitment can solve the problem for SAMP-SB, it may introduce other undesirable characteristics.

When agents bid according to the strategic policies we describe, the combinatorial protocol does not fare as well with Type 1 and Type 2 suboptimality as does SAMP-SB, with non-strategic bidding policies. The one exception for Type 2 suboptimality was NO-EQUILIB conditional on cost structures such that no price equilibrium exists. We found that 95% of these instances exhibited Type 2 suboptimality with SAMP-SB. It is not surprising that SAMP-SB, being a separately-priced-goods protocol, would exhibit increased Type 2 suboptimality, and hence lower average efficiency in networks with no price equilibrium. In contrast, the combinatorial protocol, which is not price based, did somewhat better on average in the nonequilibrium instances than it did in the price equilibrium instances of NO-EQUILIB.

7. BIDDING DEVIATIONS EXPERIMENTS

Although the base combinatorial bidding policy characterized by Equation (2) is a Bayes-Nash equilibrium for a particular network, it does not generally constitute a Bayes-Nash equilibrium for task dependency networks. Using the instances from the experiments described in Section 6, we found empirically that for each network (other than the pure parallel network in Figure 4) a single producer π could expect to obtain higher surplus by unilaterally

Network	Base combinatorial (strategic) % of instances			SAMP-SB (non-strategic) % of instances				SAMP-SB-D (non-strategic) % of instances			
	Zero	Sub	Opt	Neg	Zero	Sub	Opt	Neg	Zero	Sub	Opt
SIMPLE	32.5	1.9	65.6	0.0	0.3	0.0	99.7	0.0	0.3	0.0	99.7
UNBALANCED	48.0	3.0	49.0	5.0	0.0	7.0	88.0	0.0	4.0	1.0	95.0
NO-EQUILIB, case:											
– equilibrium exists	40.0	11.0	49.0	11.0	0.0	6.0	83.0	0.0	0.0	3.0	97.0
– no equilibrium exists	12.0	80.0	8.0	18.0	0.0	78.0	4.0	0.0	1.0	95.0	4.0
BIG	24.0	2.0	74.0	0.0	0.0	4.0	96.0	0.0	0.0	0.0	100.0
MANY-CONS	24.0	58.0	18.0	27.0	0.0	56.0	17.0	0.0	0.0	2.0	98.0

Table 1: Distribution of efficiency classes from the three protocols considered. Efficiency classes: Negative (Neg), Zero, Suboptimal (Sub), and Optimal (Opt). Note that the combinatorial protocol does not produce negative efficiency.

Network	Base combinatorial (strategic)	SAMP-SB (non-strategic)	SAMP-SB-D (non-strategic)
SIMPLE	0.799	1.000	1.000
UNBALANCED	0.519	0.747	0.960
NO-EQUILIB, case:			
– equilibrium exists	0.570	0.733	0.986
– no equilibrium exists	0.706	0.268	0.686
BIG	0.760	1.000	1.000
MANY-CONS	0.600	0.120	0.996

Table 2: Average efficiency in each network for the three protocols.

Network	Base combinatorial (strategic)	SAMP-SB (non-strategic)	SAMP-SB-D (non-strategic)
SIMPLE	0.472	0.000	0.000
UNBALANCED	0.374	−0.156	0.057
NO-EQUILIB, case:			
– equilibrium exists	0.466	0.210	0.464
– no equilibrium exists	0.521	0.137	0.555
BIG	0.456	0.001	0.001
MANY-CONS	0.447	−0.517	0.359

Table 3: Average fraction of available surplus obtained by producers in each network for the three protocols.

bidding $\max(\alpha_\pi r_\pi, \kappa_\pi)$, for some $\alpha_\pi \neq 1$, assuming that all other producers bid according to the base bidding policy. For instance, a producer in Network SIMPLE can improve its expected surplus by more than 15% with a well-chosen $\alpha_\pi > 1$.

We explored the possibility of the existence of approximate, constrained strategic equilibria we call *quasi-equilibria*. In particular, we considered quasi-equilibria such that, given a set of coefficients $\{\alpha_\pi\}$, if all producers π bid $\max(\alpha_\pi r_\pi, \kappa_\pi)$, then no producer π could improve its surplus by more than some threshold percentage $I > 0$ by bidding $\max(\alpha'_\pi r_\pi, \kappa_\pi)$ for any α'_π .

Using the base bids \tilde{r}_π computed by Monte Carlo simulation in the experiments described in Section 6, we performed a search for coefficients constituting a quasi-equilibrium for Network SIMPLE. Noting that all producers in the network have symmetric strategic problems, we limited our search to homogeneous coefficients, $\alpha = \alpha_\pi$ for all $\pi \in \Pi$. We searched for α in discrete units of .01. For a tentative $\hat{\alpha}$, we searched for deviating coefficients that varied from $\hat{\alpha}$ in discrete increments of 1%.

We found that with $\alpha = 1.13$, no producer in Network SIMPLE could improve its surplus by more than $I = 1\%$ in the random instances. No other homogeneous coefficients provided as small an

advantage to deviation.

Because producers bid higher in the quasi-equilibrium in Network SIMPLE, we find that they overbid (that is, they bid too high to win, when lower bids would have resulted in positive surplus) in more instances than they do with the base policy. Thus, as shown in Figure 16, we see that the Zero and Zero/Suboptimal/Optimal regions are bigger than with the base protocol (Figure 10). The higher bids reduce the average efficiency, as shown in Table 4. Although increased overbidding causes producers to sometimes lose out on surplus they would have obtained with lower bids, producers gain more surplus on the bids they do win. Table 4 shows that on average, producers lose surplus by bidding the quasi-equilibrium policy, as compared to the base policy.

Although we have not computed Bayes-Nash equilibria or quasi-equilibria for the other networks, we have performed limited deviation analysis with those networks. For networks BIG, MANY-CONS, and UNBALANCED, producers would expect to gain surplus by unilaterally bidding $\max(1.10r_\pi, \kappa_\pi)$, assuming all other agents use the base policy. We thus conjecture that, for these networks in a Bayes-Nash equilibrium or quasi-equilibrium, producers would generally bid higher than with the base policy. It would follow that




Figure 16: Results of quasi-equilibrium policy ($\alpha = 1.13$, $I = 1\%$) for Network SIMPLE. Number of instances: 3220.

the results for the base policy from Section 6.3 constitute an upper bound on efficiency.

Measures for Network SIMPLE	Base combinatorial	Quasi-equilibrium
Average efficiency	0.673	0.463
Producers' average fraction of available surplus	0.450	0.346

Table 4: Summary comparison of the base combinatorial bidding policy and quasi-equilibrium policy ($\alpha = 1.13$, $I = 1\%$) for Network SIMPLE.

The deviation incentives are less clear-cut in Network NO-EQUILIB. In price equilibrium instances, producers a_1 and a_3 would each gain by unilaterally bidding $\max(1.10r_\pi, \kappa_\pi)$, but other agents would not gain by such deviations. In contrast, in the non-equilibrium instances, some agents would benefit by unilaterally deviating up and some down. These results suggest that a quasi-equilibrium, if one exists, would have a more complex form than with Network SIMPLE.

8. CONCLUSIONS AND FURTHER WORK

In previous work we investigated separate, distributed markets for the problem of supply chain formation. Although suboptimal and even undesirable allocations can result, many market environments support only such negotiations. However, in some environments it is possible for a “market maker” to support a single mechanism that coordinates all auction activities in a market. A combinatorial auction, by directly linking the negotiations for all goods, ameliorates some of the coordination problems that occur in protocols involving separate negotiations.

The protocol for supply chain formation investigated here comprises a one-shot combinatorial auction mechanism along with strategic bidding policies, computed by Monte Carlo simulations. The combinatorial protocol avoids the difficulty of coordinating the acquisition of multiple inputs with the providing of its output, characteristic of independent negotiations. If agents bid non-strategically, the auction computes optimal allocations, but producers can obtain significant positive surplus with certain strategic bidding policies. However, when the available surplus is small relative to the consumers' values, the producers' strategic behavior may result in overbidding, preventing the supply chain from forming at all.

The strategic bidding policies we studied constitute Bayes-Nash equilibria for a pure parallel network, but agents in other network structures generally have incentives to deviate. We empirically found that, for a particular network, a variant of the policy forms an approximate, constrained strategic equilibrium.

It would be straightforward to extend the model in various ways, such as providing for consumer demand for multiple goods, alternative producer technologies, and multiple output choices, so long as the constraints can be expressed in linear terms. We expect that such extensions are often computationally tractable for problems of realistic size in most business-to-business (B2B) e-commerce markets. However, with increased options in the network, agents would need to perform more Monte Carlo sampling to obtain comparable fidelity in computing bids.

Although we studied only one homogeneous bidding policy for all producers, further work could involve studies of alternative policies, perhaps through evolutionary stability analysis. However, our present work suggests policies are unlikely to be uniformly successful in all networks and for all costs and values.

Finally, we believe that there are roles for both distributed separately-priced-goods auctions and combinatorial auctions in B2B negotiations. Independent market makers or cooperating businesses may be able to establish combinatorial auctions around core portions of an industry's market. However, it is inevitable that companies will need to negotiate with entities outside of the core of the industry, for it is infeasible computationally (as well as socially and politically) to manage a universal world-market combinatorial auction. Designing effective markets comprising both separately-priced-goods and combinatorial auctions remains an interesting problem for future work.

Acknowledgments

We thank Daniel Reeves for helping establish the connection between the bidding policy and Bayes-Nash equilibrium, and David Parkes for constructive comments. This work was supported by a NASA/Jet Propulsion Laboratory Graduate Student Researcher fellowship and DARPA grant F30602-97-1-0228 from the Information Survivability program. Fredrik Ygge thanks SITI and NUTEK for their financial support.

9. REFERENCES

- [1] A. Andersson, M. Tenhunen, and F. Ygge. Integer programming for combinatorial auction winner determination. In *Fourth International Conference on Multi-Agent Systems*, pages 39–46, 2000.
- [2] S. Bikhchandani and J. W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74:385–413, 1997.
- [3] F. Gul and E. Stacchetti. Walrasian equilibrium without complementarities. Technical report, Princeton University and University of Michigan, 1997.
- [4] D. Lehmann, L. I. O’Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. In *ACM Conference on Electronic Commerce*, pages 96–102, Denver, 1999.
- [5] J. K. MacKie-Mason and H. R. Varian. Generalized Vickrey auctions. Technical report, Dept. of Economics, Univ. of Michigan, July 1994.
- [6] R. P. McAfee and J. McMillan. Auctions and bidding. *Journal of Economic Literature*, 25:699–738, June 1987.
- [7] D. C. Parkes. *iBundle*: An efficient ascending price bundle auction. In *ACM Conference on Electronic Commerce*, pages

148–157, 1999.

- [8] D. C. Parkes and L. H. Ungar. Iterative combinatorial auctions: Theory and practice. In *Seventeenth National Conference on Artificial Intelligence*, pages 74–81, 2000.
- [9] D. C. Parkes and L. H. Ungar. Strategic manipulation in iterative auctions: Proxy agents and price-adjustments. In *Seventeenth National Conference on Artificial Intelligence*, pages 82–89, 2000.
- [10] S. J. Rassenti, V. L. Smith, and R. L. Bulfin. A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics*, 13:402–417, 1982.
- [11] M. H. Rothkopf, A. Pekeč, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44:1131–1147, 1998.
- [12] T. W. Sandholm and F. Ygge. On the gains and losses of speculation in equilibrium markets. In *Fifteenth International Joint Conference on Artificial Intelligence*, pages 632–638, Nagoya, Japan, 1997.
- [13] W. E. Walsh and M. P. Wellman. A market protocol for decentralized task allocation. In *Third International Conference on Multi-Agent Systems*, pages 325–332, 1998.
- [14] W. E. Walsh and M. P. Wellman. Efficiency and equilibrium in task allocation economies with hierarchical dependencies. In *Sixteenth International Joint Conference on Artificial Intelligence*, pages 520–526, 1999.
- [15] W. E. Walsh and M. P. Wellman. Modeling supply chain formation in multiagent systems. In *Agent Mediated Electronic Commerce II*, volume 1788 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2000.
- [16] M. P. Wellman and J. Hu. Conjectural equilibrium in multiagent learning. *Machine Learning*, 33:179–200, 1998.
- [17] P. R. Wurman. *Market Structure and Multidimensional Auction Design for Computational Economies*. PhD thesis, University of Michigan, 1999.
- [18] P. R. Wurman and M. P. Wellman. AkBA: A progressive, anonymous-price combinatorial auction. In *Second ACM Conference on Electronic Commerce*, 2000.
- [19] F. Ygge. *Market-Oriented Programming and its Application to Power Load Management*. PhD thesis, Lund University, 1998.