Combinatorial Auctions for Supply Chain Formation*

William E. Walsh
IBM T. J. Watson Research Center
19 Skyline Drive, Hawthorne, NY 10532 USA
wwalsh1@us.ibm.com

Michael P. Wellman
University of Michigan AI Laboratory
1101 Beal Ave, Ann Arbor, MI 48109-2110 USA
wellman@umich.edu

Fredrik Ygge
Trade Extensions
Krokebacksg. 11, 426 74 Västra Frölunda Sweden
ygge@tradeextensions.com

DRAFT `$Date: 2002/05/30 04:07:33 $`

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. Optimizing with respect to offers results in optimal allocations if the offers reflect true values and costs. But autonomous self-interested agents have an incentive to bid strategically in an attempt to gain

*Parts of this manuscript appear in a paper presented at the Second ACM Conference on Electronic Commerce (EC-00) [38].

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. We also construct Bayes-Nash equilibrium strategies for special classes of networks.

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 assembling complex production and exchange relationships between autonomous, self-interested agents

Supply chain formation can be particularly challenging when firms must contend for scarce resources at multiple levels in the supply chain structure [37]. 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.

Currently, companies address these problems by maintaining extra inventory or establishing long-term contracts based on predictions of their needs and market demand. The former is anathema to modern, no-inventory operations such as just-in-time scheduling. The latter requires significant and persistent coordination between partners in a supply chain in order to moderate uncertainty and risk. This type of coordination relies on extensive mutual knowledge between partners, which in term demands large commitments of time and trust to establish the foundations of mutual benefit. But with the rapid pace of market change and increased competitiveness, time is in short supply and alignments of interests are fleeting. Instead, fundamental changes in organizational structures may be necessary to address growing needs for responsiveness and agility. Some [6, 19] envision a not-too-distant future in which persistent, large corporations will be supplanted by temporary, agile organizations formed dynamically for ad hoc purposes. Simultaneous negotiation across the supply chain are essential to establish the viability of such "virtual corporations".

In prior work [33, 34, 35], 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 [23, 27, 41]. 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 [28] (as it is specifically in the case of supply chain formation [36]), Anderson et al. [1] have shown that a commercial mixed-integer-linear programming package can quickly solve many large allocation problems.

In the next section, we describe our model of supply chain structures, and in Section 3, review our previous work in market-based supply chain formation. We present a combinatorial auction protocol, consisting of a one-shot combinatorial auction and strategic bidding policy, 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 perform a more in-depth strategic analysis for certain classes of networks. We discuss related work in Section 8, and conclude and discuss future work in Section 9.

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 value [37]. 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 that 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 ***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. A producer's inputs and outputs are ***complementary*** in that it must acquire *each* of its inputs in order to feasibly produce its output, and the inputs have no innate value to the producer beyond their use in production. 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.

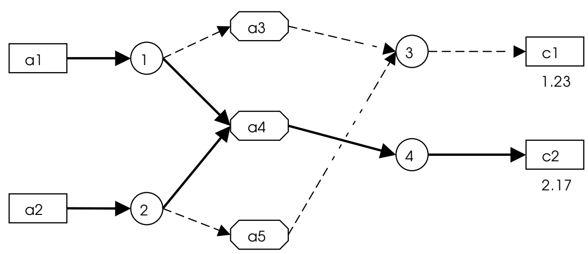

Figure 1: Task Dependency Network TWO-CONS. Circles represent goods, rectangles and octagons represent agents, and arrows represent use or provision of goods. Consumer values are indicated by amounts below their respective nodes.

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.

We analyze the quality of a solution in terms of efficiency, the ratio of the global value obtained to the maximum global value, as follows:

Definition 1 (value of an allocation) *The* value of allocation (V', E') *is:*

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

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

Definition 3 (efficiency) *The **efficiency** of allocation* (V', E') *is the fraction of the efficient value obtained by* (V', E')*, that is* $value((V', E'))/value((V^*, E^*))$.

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).

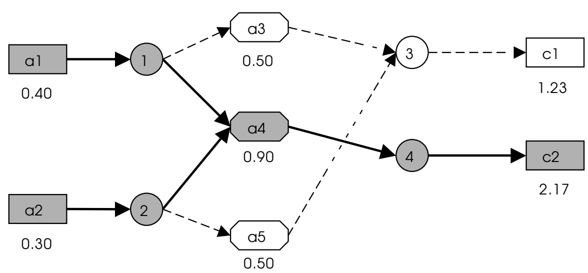

Figure 2: An efficient solution for Network TWO-CONS, given costs as shown. Shaded goods and agents and solid edges are in the solution.

In the following sections, we discuss market protocols to guide self-interested agents in finding allocations. We assume that agents wish to maximize the ***surplus*** value they obtain from the protocols. For a consumer, this is the difference between its value for the allocation and any monetary payments it makes. For a producer, this is the difference between the payments it receives and the production costs it incurs. Furthermore, we assume that the producers' costs are private information. A producer knows the distribution, but not the exact costs, of other producers.

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. We first consider separate pricing of goods because it is the default approach to market interactions. However, as we describe below, separately-priced-goods approaches are fundamentally subject to miscoordination problems in the presence of complementarities. This limitation motivates our consideration of a combinatorial auction approach to supply chain formation.

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. Specifically in our model, the conditions for competitive equilibrium are specified by the following constraints on prices for a feasible allocation: 1) a producer in the solution obtains nonnegative surplus by being active, 2) a producer not in the solution would obtain nonpositive surplus by being active, 3) a consumer in the solution obtains the good that gives it maximum positive surplus, and its total payments equal the price of that good, and 4) a consumer not in the solution would obtain nonpositive surplus from any good.

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 [3, 11, 43], and in task dependency networks in particular [35, 33]. 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. 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 [35, 33]. 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 [34, 33].

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

Zero efficiency: Failure to find a solution when an efficient solution exists. This can occur when there is insufficient slack between consumer values and producer costs.

Inferior: Forming a solution with value inferior to an efficient solution.

Dead ends: Inactive producers acquiring inputs. The paths of production lead to "dead ends" at these producers.

While other categorizations could also be useful, we believe these distinctions are qualitatively significant. Since producers at dead ends generally obtain negative surplus, they would be hesitant to participate in the protocol in the first place. If a protocol regularly produces zero efficiency, agents would consider it more worthwhile to participate in other forms of negotiation. In contrast, inferior allocations are likely to be less problematic for agents, particularly given that they cannot even be detected without perfect global knowledge of the costs and values.

Experiments on random networks [35] as well as fixed networks with random costs [33] suggest that dead ends are the dominant type of suboptimality in SAMP-SB, although zero and inferior suboptimality can and do occur.[1] Figure 3 shows the results of a run of SAMP-SB on the TWO-CONS network of Figure 2, producing a suboptimal allocation. The allocation exhibits inferior 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 dead-end suboptimality.

One way to eliminate dead-end suboptimality is to allow ***decommitment***. We define SAMP-SB-D to be the protocol SAMP-SB, followed by a phase where inactive producers ***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

[1]For comparison with the combinatorial protocol we describe below, we include the fixed-network results in Section 6.

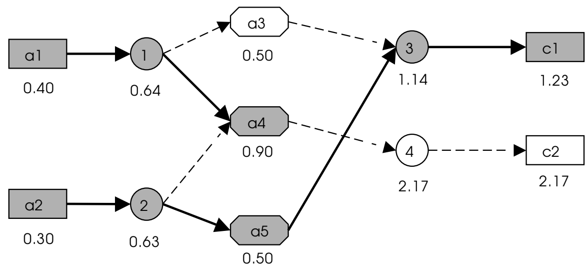

Figure 3: A suboptimal allocation generated by SAMP-SB in a network without a price equilibrium. Shaded agents and goods and solid lines are in the allocation.

that agents are guaranteed to be no worse off by participating in SAMP-SB-D. However, this is achieved by making the auction allocations nonbinding, which is undesirable to the producers who lose their output sales to decommitments. Moreover, making offers nonbinding significantly alters the strategic situation, perhaps making it less likely that agents would follow SAMP-SB in the first place. Allowing decommitment also begs the question of how to enforce the requirement that inactive producers be the only agents that decommit.

Complementarities across multiple inputs are necessary to obtain dead-end suboptimality in SAMP-SB [33], although input complementarities do not necessarily lead to dead-end suboptimality. Nevertheless, it has been widely recognized that, in both production and non-production economies with complementarities, separately-priced-goods approaches are susceptible to miscoordination. That is, when an agent values a set of goods higher than the sum of its individual value for the goods, it faces the risk of acquiring only a subset of the goods at a greater price than that subset is worth alone.

Although suboptimal and even undesirable (to some producers) allocations can result from supply chain formation, 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 dead-ends directly, though it does not inherently eliminate the other categories of sub-optimality mentioned above.

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

To simplify exposition and focus empirical study of the combinatorial protocol, we assume that a consumer c desires only a single good g_c, and denote $v_c = v_c(g_c)$. However, the framework can be easily generalized to include more complex consumer preferences.

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 reports 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

[2] Although it is possible to define "Generalized Vickrey" payments that induce truth revelation as a dominant strategy [18], 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. [16]—may present opportunities for mechanism designers to exploit.

a places a bid b_a of the form

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

where q_a^i is the integer quantity that agent a demands (positive for inputs and negative for outputs) for good g_i, and r_a is its reported willingness to pay (or be paid, in the case of negative values) 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:

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

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 by the bids*.

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

We can see immediately that if agents behave non-strategically (i.e., report their true valuations) in the combinatorial auction mechanism, then the result will always be an efficient allocation.[3] Furthermore, a producer is certain to always obtain zero surplus if it reports its true cost in the combinatorial auction. Hence, 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, as can occur in distributed protocols, such as SAMP-SB, agents may be more willing to place aggressive, strategic bids than they might when faced with separate negotiations.

[3]The 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 [25, 42], 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.

To structure the analysis, we assume that it is common knowledge that the consumers report 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 $[0,1]$. Auction theorists often like to assume that agents play Bayes-Nash equilibrium strategies [9] based on common knowledge of cost and value distributions. A set of strategies constitutes a Bayes-Nash equilibrium if no agent can improve its expected utility with respect to these distributions by unilaterally deviating from its equilibrium strategy.

Consider the degenerate (no task dependency) network, PARALLEL, of N parallel producers and one consumer with unit value, shown in Figure 4. We can derive a Bayes-Nash equilibrium for the producers in this network by recognizing its correspondence to the first-price sealed-bid auction. McAfee and McMillan [21] 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 the unique, symmetric 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

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

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.

Unfortunately, deriving such Bayes-Nash equilibria is difficult for even slightly more complicated networks. In Section 7 we identify two equilibria for one other simple network. Although one of these equilibria can be generalized to a broader

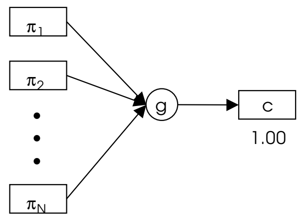

Figure 4: Network PARALLEL: A network with a closed-form Bayes-Nash equilibrium.

class of networks, this generalization is not an equilibrium for the other networks we consider. Because we cannot produce closed-form Bayes-Nash equilibrium expressions for most 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 Network PARALLEL, 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 reports

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

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

For Network PARALLEL 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 ith order statistic (ith lowest) of producer costs, this reduces to

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

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 have empirically shown that the bidding policy is not generally a Bayes-Nash equilibrium. In Section 7.2, 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” for a particular network. In Section 7.1, we present some further Bayes-Nash equilibrium analysis for a different network.

For Network PARALLEL, the combinatorial protocol we have described always chooses the lowest cost producer, thus producing efficient allocations. The winning producer always obtains a $1/N$ fraction of the available surplus. However, deriving analytic expressions of the performance of the combinatorial protocol appears to be infeasible for even slightly more complicated networks. Hence, in Section 6 we describe an empirical study of the performance of the combinatorial protocol, and compare the results with those from SAMP-SB.

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:

[4] This is not surprising, given that the Revenue Equivalence Theorem [21] implies an analogous result for the first-price sealed-bid auction.

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. We have categorized the protocols we consider along these dimensions.

	Combinatorial	Distributed
Non-strategic	one-shot combinatorial/ true costs	SAMP-SB, SAMP-SB-D
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.

It should be clear why we categorized the two representative combinatorial policies along the strategic dimension as shown. A producer reporting its true costs does not engage in any reasoning about its effect on the results, while Equation (2) directly includes the producer's contribution to the value of the efficient allocation. The proper classification of SAMP-SB and SAMP-SB-D may be somewhat less clear-cut. We have justified a producer's policy for its input offers with the argument that a producer would only gradually increase these offers to avoid overcommitting itself to inputs that it may not use in quiescence (i.e., to reduce the chance of being stuck in a dead end). Because this takes into account the possibility that the producer's actions could result in undesirable consequences, we could consider this a simple form of strategic reasoning. However, this strategic reasoning is really quite limited, for it does not account for the producer's effects throughout the network in the context of the broader bidding dynamics. In contrast, the strategic combinatorial bidding policy considers the producer's effects throughout the entire network. So we argue that SAMP-SB and SAMP-SB-D lie much closer to the non-strategic end of the spectrum than the combinatorial policy, hence categorize it as (relatively) non-strategic.

As noted in Section 4.2, 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, described in Section 6, evaluate the efficiency of the allocations obtained via simulation.

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. If a producer reports its true costs in the combinatorial auction, then it will obtain zero surplus with certainty. Additionally, because the combinatorial bids are indivisible, a producer cannot be exposed to achieving negative surplus (as can occur in SAMP-SB), even if it bids aggressively. Hence we should assume that producers will bid strategically in the combinatorial auction. In contrast, the empirical results of Section 6 show that producers can actually obtain positive surplus by bidding non-strategically in the simultaneous ascending auctions.

We have argued the plausibility of a particular bidding policy in the combinatorial mechanism (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 [25, 26, 42].

6 Performance Evaluation

6.1 Setup

In order to gain further understanding of the effectiveness of SAMP-SB, SAMP-SB-D, and the proposed combinatorial auction protocol, we simulated the protocols on sample task dependency networks. For this investigation, we chose to focus on a small set of networks exhibiting a variety of structural properties: SIMPLE (Fig-

[5] It is well established that pursuing aggressively strategic behavior can lead to pitfalls in uncertain environments [29, 39]. Proposers of particular bidding behaviors (strategic or not) in particular environments still face the burden of demonstrating reasonableness.

ure 6), UNBALANCED (Figure 7), BIGGER (Figure 8), MANY-CONS (Figure 9), and GREEDY-BAD (Figure 10), as well as TWO-CONS (Figure 1).

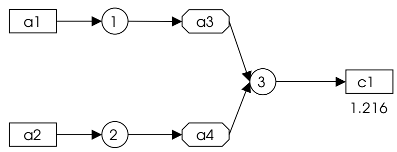

Figure 6: Network SIMPLE.

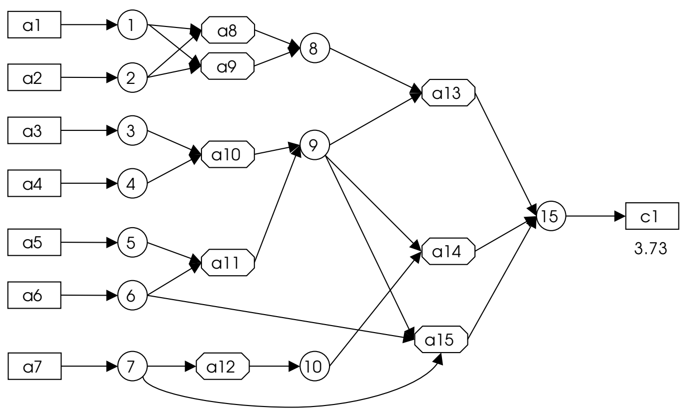

Figure 7: Network UNBALANCED.

We ran experiments on multiple instances of each network. For each instance we randomly chose producer costs uniformly from $[0,1]$, but 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. We discarded all instances whose efficient solutions had value zero.

To test the effect of competitive equilibrium existence on the performance of the protocols, we generated instances of UNBALANCED, TWO-CONS, and GREEDY-BAD with costs that admit competitive equilibrium and with costs that do not. Because SIMPLE and MANY-CONS are polytrees, all instances thereof have competitive equilibria [33]. We were not able to generate no-equilibrium instances of BIGGER with the given cost distributions.

To generate an instance with a desired type of cost structure (equilibrium or not) we repeatedly chose sets of producer costs randomly from the uniform dis-

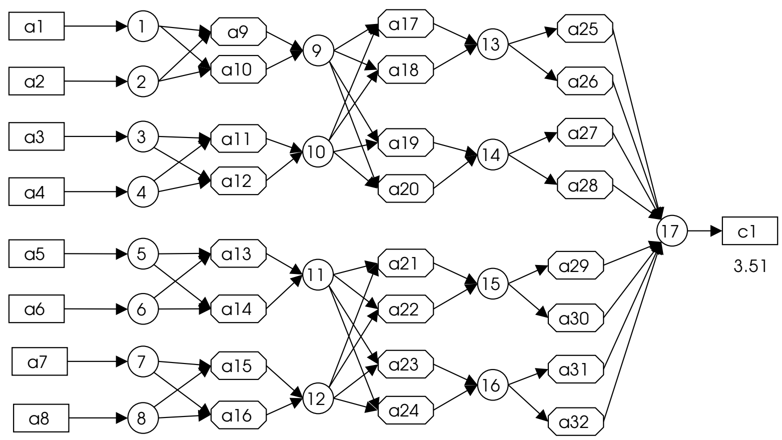

Figure 8: Network BIGGER.

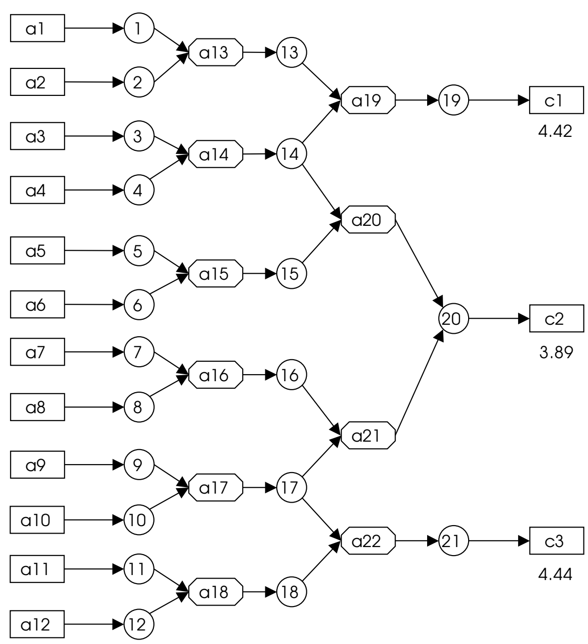

Figure 9: Network MANY-CONS.

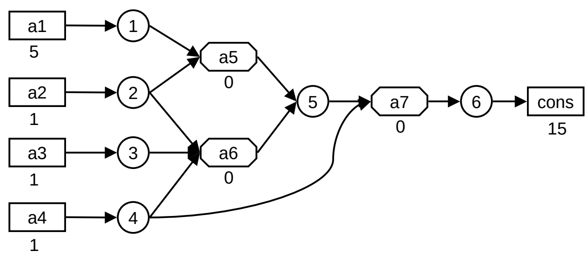

Figure 10: Network GREEDY-BAD.

tribution until the desired property was met. In the experiments, we determined whether competitive equilibrium existed—given complete information about the network structure, values, and costs—using the fact that competitive equilibria are always efficient [33], along with the following procedure. Given an optimal allocation (V^*, E^*), we attempt to solve the system of linear equations that characterize a competitive equilibrium. These can be determined straightforwardly from the verbal descriptions in Section 3 [38].

For each type of cost structure in each network, we tested 100 random instances, with the exception of SIMPLE. For this it was feasible to run many more instances, and we happened to run 3210. For each instance and each protocol, we measured the ***efficiency***—the fraction of the efficient value—attained by SAMP-SB and SAMP-SB-D. We also measured the fraction of available surplus (i.e., fraction of the value of an optimal solutions) obtained by the producers.

6.2 Implementation

Lacking a general closed form for the combinatorial bidding policy, in the experiments we employed 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 $\tilde{r}_\pi$, 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 $\tilde{r}_\pi$ values. If it does not participate in the efficient allocation for any sample, it refrains from bidding.

Determining whether a solution exists in a task dependency network is NP-complete [36], implying that computing an optimal allocation is NP-hard. Still, sophisticated software can many instances in practice [1]. We used CPLEX to compute the combinatorial auction allocation (winner determination) problem of $\Psi(B)$, the Δ_π and Δ^* values for the simulations (which are essentially equivalent to winner determination, but with different parameters), and the linear constraints on equilibrium existence. 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

In experiments run, the efficiency of the combinatorial protocol was strongly dependent on the available surplus, $\hat{\Psi}(A)$. Figures 11–19 plot the efficiency achieved by the combinatorial protocol on the 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.

For all networks, the combinatorial protocol efficiency was always zero for some region in the low range of the available surplus. For all networks, with the notable exception of those instances conditional on no equilibrium, there were "Essentially Optimal" regions (most instances optimal, with a few slightly suboptimal instances) at the high range of available surplus. Some networks also exhibited regions of mixed efficiency classes in the middle regions of available surplus.

In contrast to the combinatorial protocol, we do not observe a strong relationship between the available surplus and the efficiency classes for SAMP-SB or SAMP-SB-D. Nevertheless, we do observe distinct and varied efficiency classes. Tables 1 and 2 shows the percentage of instances in each class, for the three protocols and each network.

Recall (from Section 3) that efficiency loss in SAMP-SB can be attributable to any of three causes: dead ends, failure to form a solution when a positive-valued solution exists, and finding an inferior solution. We can infer the percentage of instances exhibiting dead-end suboptimality in SAMP-SB by examining the differences between SAMP-SB-D and SAMP-SB totaled over the Negative, Zero,

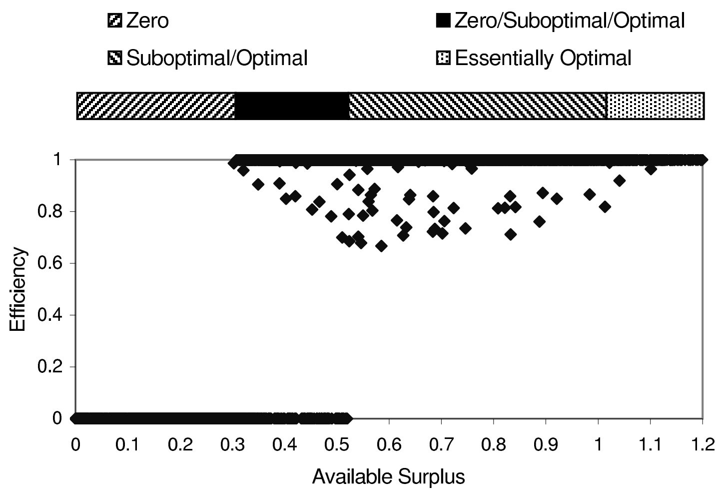

Figure 11: Results of combinatorial protocol experiments for Network SIMPLE. Number of instances: 3210.

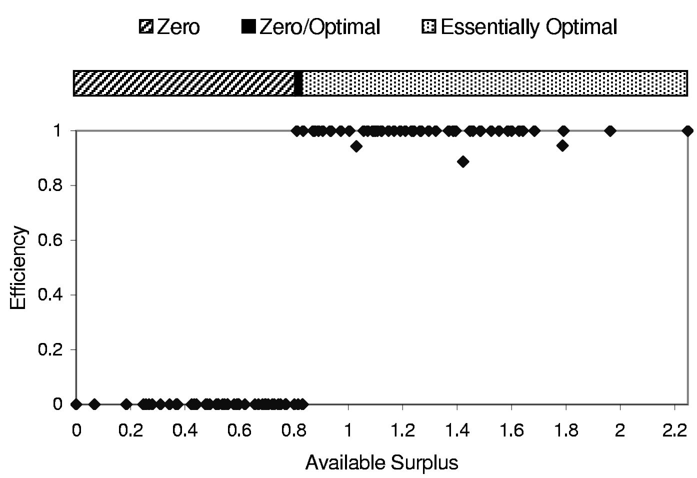

Figure 12: Results of combinatorial protocol experiments for Network UNBALANCED conditional on cost structures for which competitive equilibrium exists. Number of instances: 100.

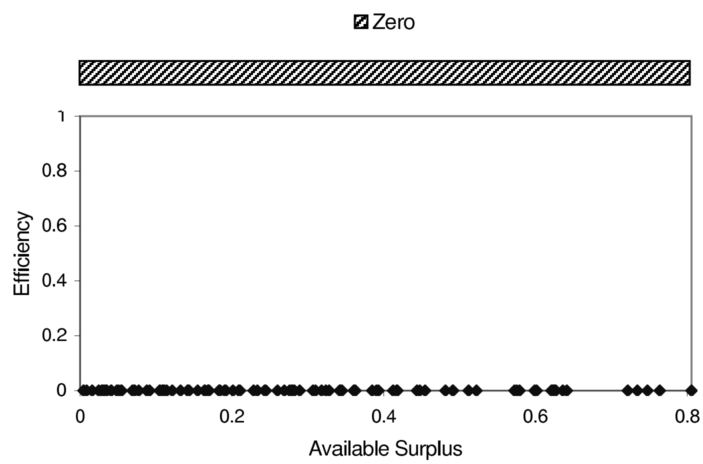

Figure 13: Results of combinatorial protocol experiments for Network UNBALANCED conditional on cost structures for which no competitive equilibrium exists. Number of instances: 100.

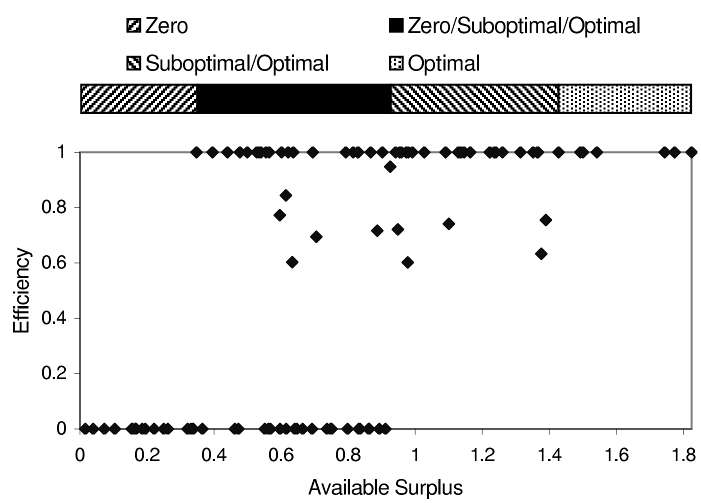

Figure 14: Results of combinatorial protocol experiments for Network TWO-CONS conditional on cost structures for which competitive equilibrium exists. Number of instances: 100.

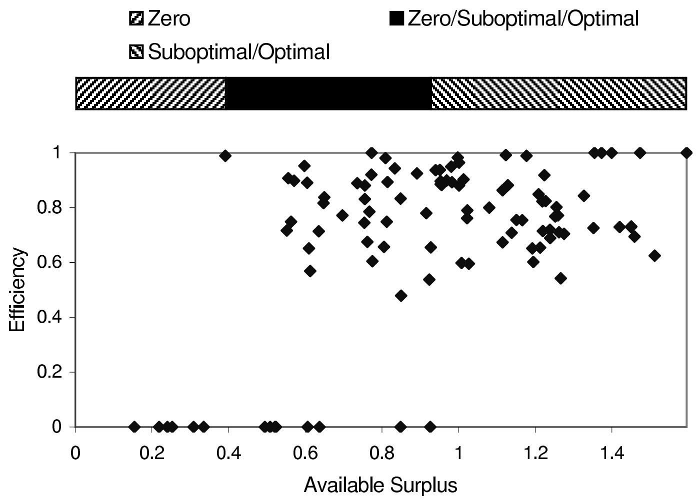

Figure 15: Results of combinatorial protocol experiments for Network TWO-CONS conditional on cost structures for which no competitive equilibrium exists. Number of instances: 100.

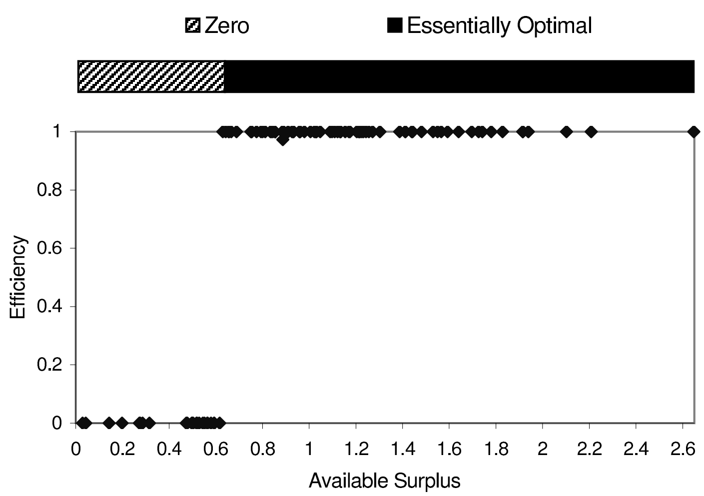

Figure 16: Results of combinatorial protocol experiments for Network BIGGER. Number of instances: 100.

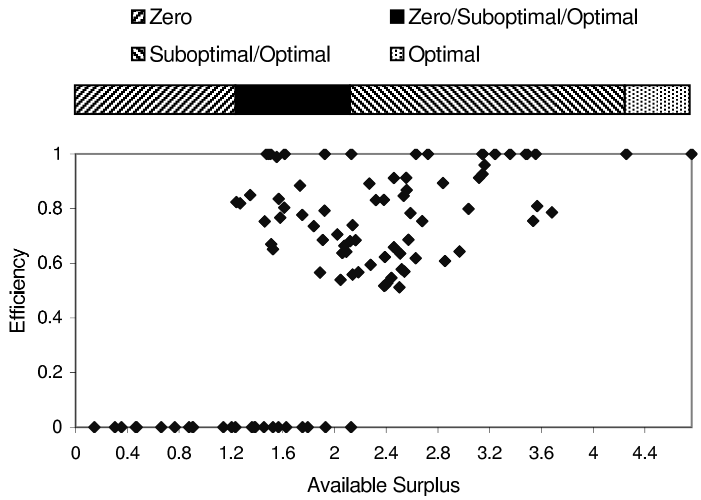

Figure 17: Results of combinatorial protocol experiments for Network MANY-CONS. Number of instances: 100.

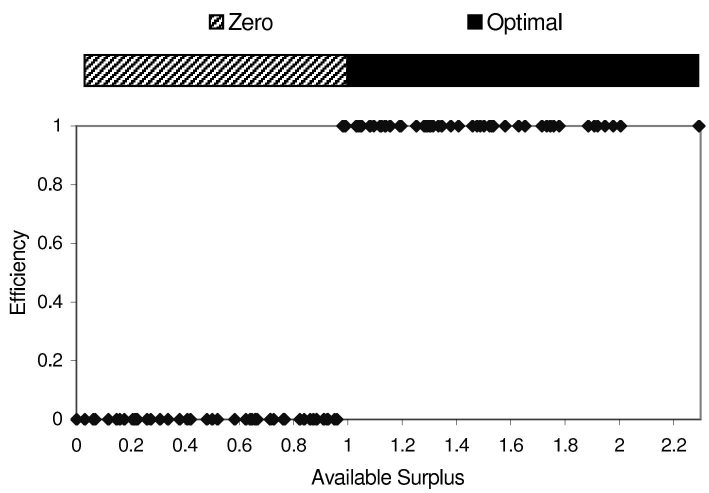

Figure 18: Results of combinatorial protocol experiments for Network GREEDY-BAD, conditional on cost structures for which equilibrium exists. Number of instances: 100.

Network	Combinatorial (strategic) % of instances Zero	Sub	Opt
SIMPLE	32.4	2.0	65.6
UNBALANCED, case:			
– equilibrium exists	47.0	3.0	50.0
– no equilibrium exists	100.0	0.0	0.0
TWO-CONS, case:			
– equilibrium exists	40.0	11.0	49.0
– no equilibrium exists	14.0	80.0	6.0
BIGGER	24.0	2.0	74.0
MANY-CONS	24.0	58.0	18.0
GREEDY-BAD, case:			
– equilibrium exists	47.0	0.0	53.0
– no equilibrium exists	99.0	0.0	1.0

Table 1: Distribution of efficiency classes from the combinatorial protocol. Efficiency classes: Zero, Suboptimal (Sub), and Optimal (Opt).

and Suboptimal columns in Table 2. Decommitment does not affect the contribution of no-solution and suboptimal-solution suboptimality, but helps reveal them by eliminating dead-end suboptimality. Hence, we can infer the percentage of instances exhibiting no-solution and suboptimal-solution suboptimality in SAMP-SB by examining the Zero and Suboptimal columns of SAMP-SB-D, respectively.

The combinatorial protocol solves the problem of dead-end suboptimality by design, which can produce significant benefits in some networks, as compared to SAMP-SB. Although decommitment solves the problem for SAMP-SB, recall that it introduces some other problems.

When agents bid according to the strategic policies we describe, the combinatorial protocol produces more instances of no-solution and suboptimal-solution suboptimality than does SAMP-SB (which has non-strategic bidding policies). Observe that only in Networks TWO-CONS and MANY-CONS does the combinatorial protocol produce a sizeable fraction of suboptimal but nonzero instances.[6]

[6]Although Figure 11 seems to suggest that there were also a significant number of suboptimal

Network	SAMP-SB (non-strategic) % of instances				SAMP-SB-D (non-strategic) % of instances		
	Neg	Zero	Sub	Opt	Zero	Sub	Opt
SIMPLE	0.0	0.3	0.0	99.7	0.3	0.0	99.7
UNBALANCED, case:							
– equilibrium exists	5.0	1.0	7.0	87.0	1.0	1.0	98.0
– no equilibrium exists	100.0	0.0	0.0	0.0	100.0	0.0	0.0
TWO-CONS, case:							
– equilibrium exists	11.0	0.0	6.0	83.0	0.0	3.0	97.0
– no equilibrium exists	18.0	0.0	78.0	4.0	1.0	95.0	4.0
BIGGER	0.0	0.0	4.0	96.0	0.0	0.0	100.0
MANY-CONS	27.0	0.0	56.0	17.0	0.0	2.0	98.0
GREEDY-BAD, case:							
– equilibrium exists	4.0	0.0	21.0	75.0	1.0	0.0	99.0
– no equilibrium exists	100.0	0.0	0.0	0.0	100.0	0.0	0.0

Table 2: Distribution of efficiency classes from SAMP-SB and SAMP-SB-D. Efficiency classes: Negative (Neg), Zero, Suboptimal (Sub), and Optimal (Opt).

We can explain this phenomenon by the fact that these are the only two networks with multiple consumers. In a network with a single consumer, if the producers overbid (that is, they report too high to win, when lower reports would have resulted in positive surplus), then the combinatorial auction will generate no solution, resulting in zero efficiency. But if there are multiple consumers desiring different goods, then even if producers in the efficient solution overbid, there is still the possibility that producers' offers do not exceed the value of a consumer in a suboptimal solution. In this case, we would get a suboptimal, nonzero solution from the combinatorial protocol.

A curious result is that the combinatorial protocol produces a much greater fraction of suboptimal allocations in the no-equilibrium instances than in the equilibrium instances of Network TWO-CONS. This is not a general pattern, but is the result of specific conditions arising from the cost structures in no-equilibrium instances of this network. It can be shown [33] that the constraints on consumer c2

instances in Network SIMPLE, they actually amount to only 2%, as shown in Table 1. There is a greater *number* of suboptimal instances because the sample is larger.

being in an efficient allocation, combined with a necessary condition for nonexistence of competitive equilibrium in this network result in producer a4 overcomputing Δ^* for no-equilibrium instances (note that the agents do not know whether they are in a no-equilibrium instance). As a result, a4 tends to overbid in no-equilibrium instances (when it is in the efficient allocation), allowing the second-best solution to be chosen by the auction, thus giving a suboptimal but positive-value allocation.

Network	Combinatorial (strategic)	SAMP-SB (non-strategic)	SAMP-SB-D (non-strategic)
SIMPLE	0.673	0.997	0.997
UNBALANCED			
– equilibrium exists	0.528	0.867	0.990
– no equilibrium exists	0.000	−20.08	0.000
TWO-CONS, case:			
– equilibrium exists	0.570	0.733	0.986
– no equilibrium exists	0.692	0.268	0.686
BIGGER	0.760	1.000	1.000
MANY-CONS	0.601	0.120	0.996
GREEDY-BAD, case:			
– equilibrium exists:	0.530	−5.32	0.990
– no equilibrium exists:	0.010	−18.23	0.000

Table 3: Average efficiency in each network for the three protocols.

Table 3 shows the average efficiency attained by the three protocols, factored by network and equilibrium existence (where relevant). We see, from the difference between the SAMP-SB-D and SAMP-SB columns, that dead ends are a significant source of inefficiency. Additionally, existence of competitive equilibrium has a significant effect on the performance of the protocols. In these networks, SAMP-SB-D produces nearly perfect efficiency when competitive equilibrium exists (recall that all studied instances of SIMPLE, BIGGER, and MANY-CONS have equilibria), but is is much less effective when equilibrium does not exist, in fact failing to find any solutions in the no-equilibrium cases of UNBALANCED and GREEDY-CONS.

It is not surprising that SAMP-SB, being a distributed, price-based protocol would produce higher efficiency in instances that admit competitive equilibrium than in instances of the same network that do not admit equilibrium. However, it is less clear that the combinatorial protocol should be sensitive to the existence of competitive equilibrium. Moreover, we do not observe consistency across networks in which type of cost structures result in better efficiency for the combinatorial protocol. With Network TWO-CONS the combinatorial protocol does somewhat better when equilibrium does *not* exist, while for the other networks it does significantly better when equilibrium *does* exist.

Network	Combinatorial (strategic)	SAMP-SB (non-strategic)	SAMP-SB-D (non-strategic)
UNBALANCED	6.52×10^{-18}	6.27×10^{-30}	8.23×10^{-101}
TWO-CONS	3.26×10^{-2}	5.15×10^{-7}	1.43×10^{-22}
GREEDY-BAD	2.17×10^{-17}	1.41×10^{-1}	8.04×10^{-101}

Table 4: P-values computed with the Student's t-Test. The t-Test compared the efficiency means of instances that admit competitive equilibrium and those that do not admit competitive equilibrium.

To check whether these differences in performance are significant, we performed Student's t-Tests for each protocol, comparing the mean efficiencies of instances that admit competitive equilibrium with the means of those instances that do not admit competitive equilibrium. Table 4 shows the results of the tests, indicating the p-values that the means of equilibrium and no-equilibrium instances came from the same underlying population. Only two entries have positive p-values within three decimal places. Still the combinatorial/TWO-CONS p-value is small enough to reject the hypothesis that the equilibrium and no-equilibrium means are actually the same. On the face of it, the high SAMP-SB/GREEDY-BAD p-value suggests that we cannot safely reject the hypothesis. However, inspection of the data indicates that this high probability results from one outlying equilibrium instance with a large negative efficiency. Indeed, the fact that SAMP-SB-D/GREEDY-BAD entry is zero gives us further evidence that it is unlikely that the equilibrium and no-equilibrium means are the same for SAMP-SB in Network GREEDY-BAD.

We must be cautious with these t-Test results because the distribution of mean efficiencies, being neither uniform nor even unimodal, does not meet the the formal criteria for the applicability of a t-Test. However, an inspection of the data supports the conclusions we reach from the t-Tests. In the instances of Networks UNBALANCED and GREEDY-BAD that do not admit competitive equilibrium, SAMP-SB-D and the combinatorial protocol essentially always produce zero efficiency, but produce perfect efficiency in many of the instances that do admit competitive equilibrium. With Network TWO-CONS, SAMP-SB-D always produces essentially optimal results in instances that admit competitive equilibrium, but mostly produces suboptimal results in the instances that do not admit competitive equilibrium. With Network TWO-CONS, the combinatorial protocol produces mostly optimal or zero efficiency results in the instances that admit competitive equilibrium, but produces a cloud of suboptimal, but positive results for the no-equilibrium instances. Thus we can safely conclude that the mean efficiencies really are different for equilibrium and no-equilibrium instances.

Network	Base combinatorial (strategic)	SAMP-SB (non-strategic)	SAMP-SB-D
SIMPLE	0.472	0.000	0.000
UNBALANCED, case:			
– equilibrium exists	0.382	−0.041	0.082
– no equilibrium exists	0.000	−20.09	0.000
TWO-CONS, case:			
– equilibrium exists	0.466	0.210	0.464
– no equilibrium exists	0.518	0.137	0.555
BIGGER	0.466	0.001	0.001
MANY-CONS	0.447	−0.517	0.359
GREEDY-BAD, case:			
– equilibrium exists	0.396	−6.08	0.137
– no equilibrium exists	0.010	−18.11	0.000

Table 5: Average fraction of available surplus obtained by producers in each network for the three protocols.

Table 5 shows the average fraction of available surplus obtained by producers, respectively, in each network, for the three protocols. Recall that in the combinatorial auction, extra surplus (available as specified by the strategic bids) not taken by producers is distributed evenly among all consumers. We see that producers can obtain significant positive surplus with the strategic bidding policy in the combinatorial auction. Perhaps surprisingly, in some networks the producers can gain significant surplus with the SAMP-SB-D protocol, even though they are bidding to obtain zero surplus. The reason for this is that a producer's output offer indicates the minimum amount it is willing to accept in exchange for its output.

But rising buy offers can cause the price to rise above the producer's output offer. This could happen in cases when it is necessary to block out certain agents to have a feasible allocation in quiescence. Note however, that the decommitment step is needed to distribute high average surplus to the producers. Without decommitment, the average producer surplus can be highly negative, as shown in the SAMP-SB column.

7 Further Strategic Analysis of the Combinatorial Protocol

To this point, we have assumed a particular bidding policy that we have argued is plausible for the combinatorial auction. Although we can show that it constitutes a Bayes-Nash equilibrium for one simple class of networks (namely the class of simple, parallel networks shown in Figure 4), we can show that it is not an equilibrium for all task dependency networks. To gain greater insight into the strategic problem in the combinatorial auction, we perform further analysis for two simple networks. In Section 7.1 we describe a very simple network for which we can compute *multiple* Bayes-Nash equilibria. In Section 7.2 we explore a weaker notion of approximate, constrained strategic equilibrium in Network SIMPLE.

7.1 Special Case Bayes-Nash Equilibrium

We consider Network NASH-DEMAND, shown in Figure 20. We assume the same information, cost distributions, and combinatorial auction as before. Consumer c has value $v = v_c$ for obtaining good g_2. For convenience, we denote $b_\pi = -r_\pi$ for a producer π (recall that r_π is negative for a producer).

Given b_1 and b_2, if $b_1 + b_2 \leq v$ then all agents win their bids in the auction for perfect efficiency and the surplus of producer π is $b_\pi - \kappa_\pi$. Otherwise, all agents receive zero surplus. Thus the average efficiency obtained with any set of bidding policies is equal to the probability that the solution is computed by the combinatorial auction.

The strategic game bears some similarity to the ***asymmetric information version of the Nash demand game*** [4, 17, 30], for which continua of equilibria exist in various forms. In the asymmetric information Nash demand game, there is a buyer with value v' and a seller with cost κ', each drawn from uniform distributions that are common knowledge. The buyer and seller place bids to exchange a good

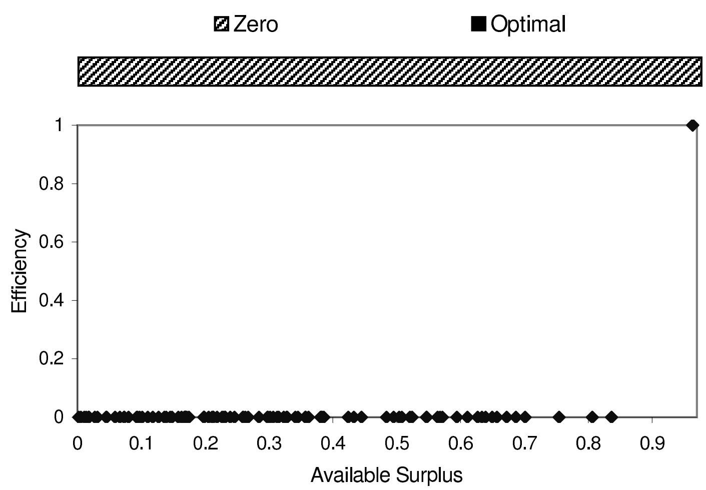

Figure 19: Results of combinatorial protocol experiments for Network GREEDY-BAD, conditional on cost structures for which no equilibrium exists. Number of instances: 100.

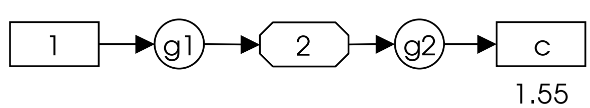

Figure 20: Network NASH-DEMAND.

from the seller to the buyer. If the buyer's report is no greater than the seller's offer, then they exchange the good at a price halfway between the two offers. In our present example, we could model the seller as producer 1, with $\kappa = \kappa_1$. Because we assume that the consumer reports its true value, which is common knowledge, we can model the buyer as the combination of the consumer c and producer 2, with $v' = v - \kappa_{\pi_2}$. However, to maintain strategic equivalence with our original game, we would have to assume that the excess surplus is "thrown away", rather than given to the consumer. This reveals a significant difference from the Nash demand game, in which the buyer and seller split the surplus difference between their reports.

7.1.1 Closed Form of the Base Combinatorial Bidding Policy

The closed form for the base combinatorial bidding policy (2) in Network NASH-DEMAND, for a producer π is:

$$b_\pi = \kappa_\pi + 1/2 \begin{cases} v - \kappa_\pi - 1/2 & \text{if } \kappa_\pi < v - 1 \\ (v - \kappa_\pi)/2 & \text{otherwise} \end{cases} \tag{4}$$

We can show that this bidding policy is not a Bayes-Nash equilibrium for this network. Depending on its cost, a producer may improve its expected surplus if it reports a lower or higher amount.

We are unable to analytically derive Bayes-Nash equilibria for Network NASH-DEMAND, as we could for Network PARALLEL. Instead, we have identified two equilibria by guessing, with inspiration from the Nash demand game literature. While the fixed-price equilibrium shown in Section 7.1.2 can be identified with a relatively little thought, identifying the piecewise-linear equilibrium shown in Section 7.1.3 requires more analysis. To identify the latter, we first guessed at the form and then computed the linear parameters and breakpoints.

7.1.2 Fixed-Price Bayes-Nash Equilibrium

We consider the following fixed-price policy for Network NASH-DEMAND, and for any $x \in [0, v]$:

$$b_1 = \begin{cases} x & \text{if } \kappa_1 \leq x \\ v & \text{otherwise} \end{cases} \tag{5}$$

$$b_2 = \begin{cases} v - x & \text{if } \kappa_2 \leq v - x \\ v & \text{otherwise} \end{cases} \tag{6}$$

Theorem 1 *Equations 5 and 6 constitute an asymmetric Bayes-Nash equilibrium for Network* NASH-DEMAND*, for any* $x \in [0, v]$.

Proof. To see that the fixed-price policy is a Bayes-Nash equilibrium, assume that producer 2 reports according to Equation 6. Since producer 1 would never improve its chance of winning with a report below x, and can never win with a report above x, producer 1 effectively has the choice of (possibly) winning with a report of x. Producer 1 would choose to win at x precisely when $\kappa_1 \leq x$. Hence, (5) is a best response by producer 1. By a similar argument, (6) is also a best response by producer 2, if producer 1 follows Equation 5. □

With this bidding policy, the producers jointly obtain 100% of the surplus iff they both win. Hence the expected efficiency is equal to the expected fraction of the available surplus obtained by the producers (recall also that the expected efficiency is equal to the probability that a solution is computed).

It turns out that the expected efficiency is maximized when $x = v/2$. If $v \geq 2$, then the agents will always win for any $x \in [1, v]$ (which includes $v/2$). If $v < 2$ then the probability that the agents win is $x(v - x)$, which is maximized at $v/2$.

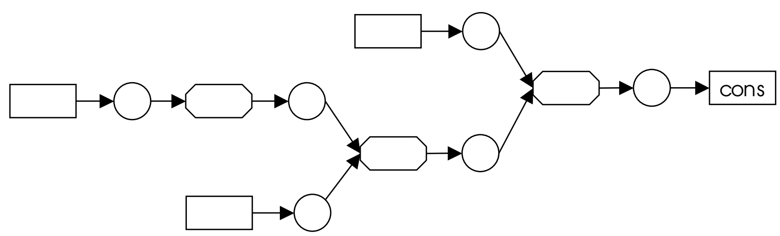

Figure 21: Network MULTI-NASH-DEMAND.

We can generalize this fixed-price policy to apply to any connected, single consumer network in which $|\{\langle a, g\rangle \in E\}| = |\{\langle g, a'\rangle \in E\}| = 1$ for all goods g (i.e., for every good there is exactly one potential buyer and one potential seller), an example of which is Network MULTI-NASH-DEMAND, shown in Figure 21. Because there is exactly one solution for any network (V, E) in this class, it is easy to see that the following is a Bayes-Nash equilibrium for (V, E):

$$b_\pi = \begin{cases} v/x_\pi & \text{if } \kappa_\pi \leq x_\pi \\ v & \text{otherwise} \end{cases} \tag{7}$$

As with Network NASH-DEMAND, which is a special case of this class, producers simply have the decision of whether or not to participate in the solution, and would choose to do so iff $\kappa_\pi \leq v/x_\pi$.

7.1.3 A Piecewise-Linear Bayes-Nash Equilibrium

Although the fixed-price policy described in the previous section is a Bayes-Nash equilibrium, it is not related to the costs of the producers. In this section we consider the following piecewise-linear policy that is more closely tied to the producers' costs:

$$b_\pi = \begin{cases} v/2 & \text{if } v > 3 \\ 2v/3 - 1/2 & \text{if } v \leq 3 \text{ and } \kappa_\pi \leq 2v/3 - 1 \\ \kappa_\pi/2 + v/3 & \text{otherwise} \end{cases} \tag{8}$$

Theorem 2 *Equation 8 constitutes a symmetric Bayes-Nash equilibrium for Network* NASH-DEMAND.

A proof of this theorem can be found in the first author's doctoral thesis [33].

7.1.4 Comparison of Bidding Policies

Policy	Efficiency	Producer Surplus
Base	0.481	0.397
Fixed	0.670	0.670
Piecewise Linear	0.594	0.490

Table 6: Efficiency and fraction of available surplus obtained by producers with three combinatorial bidding policies: Base (4), Fixed (5, 6), and Piecewise Linear (8).

Table 6 shows the expected efficiency and expected fraction of available surplus obtained by producers using three bidding policies: Base (Equation 4), Fixed price (Equations 5 and 6), and Piecewise Linear (Equation 8). The consumer value is fixed at $v = (10 - \sqrt{5})/5 = 1.55$, the value which gives a positive-value solution in 90% of the instances. For the fixed-price policy, the results are for $x = v/2$, which maximizes expected efficiency.

These results were obtained by extensive simulation[7] (millions of Monte Carlo trials—all reported sample means have 99.9% confidence intervals of ± 0.0005),

[7]We thank Daniel Reeves for performing these simulations, and also for his significant role in producing the analytic results for Network NASH-DEMAND.

with producer costs chosen randomly from the uniform distribution $[0,1]$, and considering only instances with positive available surplus. The efficiency results for the base policy were also confirmed analytically.

Observe that we obtain higher efficiency and higher producer surplus for both of the Bayes-Nash equilibria than with the base bidding policy. So for this network, either of these equilibria would be preferable to the base policy.

We would like to generalize or revise the base policy to include Bayes-Nash equilibria for a broader class of networks. Even better would be a general method for automatically computing Bayes-Nash equilibria. Either of these goals appear to require non-trivial advances. To date, we have explored only two forms of Bayes-Nash equilibrium for Network NASH-DEMAND, although the fixed-price form represents a continuum of policies, parametrized by x. The fact that we arrived at these equilibria simply by guessing at the forms and then proving their equilibrium properties, suggests that there could be a multitude of other equilibria for this network. These results then beg the question of how agents should choose *which* equilibrium to play. Indeed, equilibrium selection is an open problem [13], with many refinements for selecting equilibria proposed within the field of game theory.

7.2 Deviation and Quasi-Equilibrium

Using the instances from the experiments described in Section 6, we found empirically that for Networks SIMPLE, UNBALANCED, TWO-CONS, BIGGER, and MANY-CONS, a single Producer π could expect to obtain higher surplus by unilaterally reporting $\max(\eta_\pi r_\pi,\ \kappa_\pi)$, for some $\eta_\pi \neq 1$, assuming that all other producers report 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 $\eta_\pi > 1$.

Unable to compute a Bayes-Nash equilibrium for these networks, 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 $\{\eta_\pi\}$, if all producers π report $\max(\eta_\pi r_\pi,\ \kappa_\pi)$, then no producer π could improve its surplus by more than some threshold percentage $I > 0$ by reporting $\max(\eta'_\pi r_\pi,\ \kappa_\pi)$ for any η'_π.

Using the base reports $\tilde{r}_\pi$ 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 the search to homogeneous coefficients, $\eta = \eta_\pi$ for all $\pi \in \Pi$. We searched for η in discrete units of .01. To determine whether a given candidate $\hat{\eta}$ constitutes a quasi-equilibrium, we searched for deviating coefficients that varied from $\hat{\eta}$ in discrete increments of 1%. If this process determined that no producer would gain more than I% improvement, we concluded that $\hat{\eta}$ constitutes a homogeneous quasi-equilibrium.

For a tentative $\hat{\eta}$, we searched for deviating coefficients that varied from $\hat{\eta}$ in discrete increments of 1%.

We found that with $\eta = 1.13$, 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.

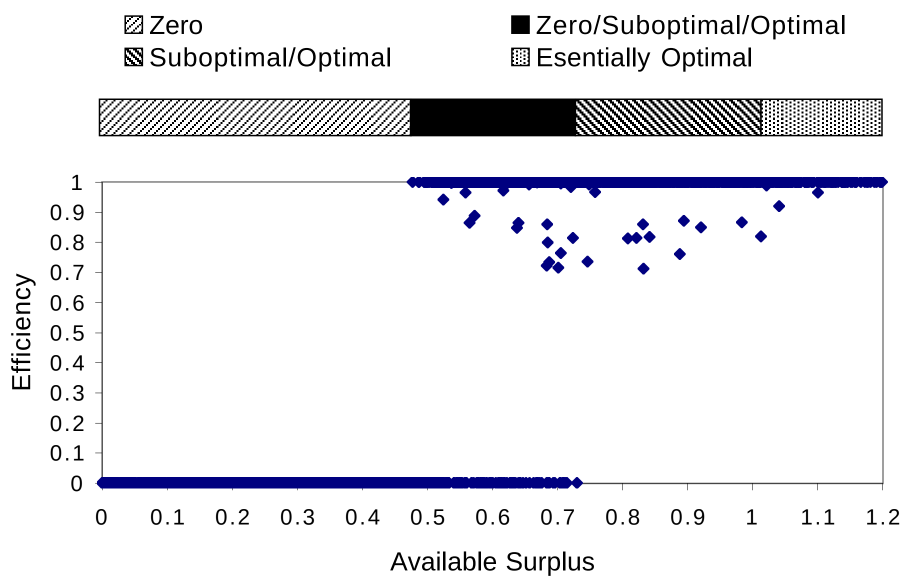

Figure 22: Results of quasi-equilibrium policy ($\eta = 1.13$, $I = 1\%$) for Network SIMPLE. Number of instances: 3220.

Because producers report higher in the quasi-equilibrium in Network SIMPLE, they often overbid in more instances than they do with the base policy. Thus, as shown in Figure 22, we see that the Zero and Zero/Suboptimal/Optimal regions are bigger than with the base protocol (Figure 11). The higher reports reduce the average efficiency, as shown in Table 7. Although increased overbidding causes producers to sometimes lose out on surplus they would have obtained with lower reports, producers gain more surplus on the reports they do win. Table 7 shows

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

that the former dominates, for on average, producers obtain lower surplus by reporting according to 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 BIGGER, MANY-CONS, and UNBALANCED (unconditional on competitive equilibrium existence), producers would expect to gain surplus by unilaterally reporting $\max(1.10 r_\pi,\ \kappa_\pi)$, assuming all other agents use the base policy. We thus conjecture that, in a Bayes-Nash or quasi-equilibrium for these networks, producers would generally report higher than with the base policy. It would follow that the results for the base policy from Section 6.3 constitute an upper bound on efficiency.

The deviation incentives are less clear-cut in Network TWO-CONS. In competitive equilibrium instances, producers a1 and a3 would each gain by unilaterally reporting $\max(1.10 r_\pi,\ \kappa_\pi)$, but other agents would not gain by such deviations. In contrast, in the no-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.

7.3 Implications for Bayes-Nash Equilibrium

Our method in this paper is to compute Bayes-Nash equilibria individually for special-case network structures, and to study in depth a plausible bidding policy that generalizes an equilibrium for a particular network structure. Ideally, we would be able to compute a closed-form equilibrium that can be generalized to all network structures in a reasonable fashion (e.g., by parametrization, but not by case-based enumeration).

Because bidding Policy (1) is the unique, symmetric equilibrium for Network PARALLEL, any candidate universal symmetric equilibrium policy must reduce to (1) for that network. While the equilibrium (1) addresses producer competition in a single-level network, a universal equilibrium must also take into account the sharing of surplus between producers active in a solution (i.e., sequences of producers and producers of complementary inputs to another producer). The generalization of Policy (2) does do so in a reasonable way, yet is not an equilibrium for other networks we studied. On the other hand, Policy (7) addresses the surplus sharing between producers, and is an equilibrium when there is no competition between producers. We have not identified an equilibrium that applies to our supply chain model universally, and our present work provides discouraging evidence on the open question of whether such an equilibrium even exists in a reasonable form.

8 Related Work

As we noted above, the miscoordination problems of separately-priced-goods approaches are widely recognized. This recognition has led to much recent attention to combinatorial auctions [7], and their application (actual or proposed) to a variety of problems, including allocating space shuttle payload priority [15], airport landing strip reservations [27], transportation services [14], factory scheduling [40], packaging material procurement [31], and the FCC spectrum auctions [8], among others.

The bulk of the combinatorial auction work has focused on single-point one-sided auctions (i.e., with a single seller and multiple buyers, or vice-versa), which cannot guarantee full coordination of an allocation across the multiple-level supply chains our task dependency network model encompasses. One exception we are aware of, besides our work, is that of Hunsberger and Grosz [12], who study the effect on combining production roles in the Shared Plans framework to speed up the computation of a combinatorial auction.

Parkes [23] and Wurman and Wellman [42] present one-sided combinatorial auctions that allow incremental revelation of preferences through iterated bidding. In our present work we assume that producers have simple production capabilities that they can reasonably represent in a single bid. In real supply chains, we would expect that some companies would have alternate means of production and choices of production goods. Our one-shot bidding model would be inadequate if the combinations of options are numerous and require complex scheduling, planning, and optimization for their determination. For these cases, iterative combinatorial auctions would be an important advance for our supply chain formation.

The difficulty of computing Bayes-Nash equilibria is a salient theme of our present work. In contrast, the Groves [10]-Clarke [5]-Vickrey [32] Auction (also called the Generalized Vickrey Auction (GVA) [18]), is incentive compatible and efficient. That is, it is a dominant strategy for all agents to report their values truthfully, and efficient allocations are obtained when agents play their dominant strategies. However, when both buyers and sellers can bid (as required for supply chain formation), the auction must generally subsidize the bidders to obtain this result [22]. Since we should generally expect that the auctioneer would not be amenable to losing money, the GVA is not a viable candidate for supply chain formation.

Alternatively, we can maintain budget balance at the auctioneer if we allow that agents to obtain negative surplus or give up perfect efficiency. Babaioff and Nisan [2] take the latter approach by applying a trade reduction rule, based on McAfee's [20] double auction rule, in a distributed fashion. The rule explicitly throws away efficient trades to obtain incentive compatibility and budget balance for linear supply chains. Parkes et al. [24] present a class of two-sided, single-level auctions that give high efficiency and low incentive to deviate from truthful bidding. We would be interested to see extensions of these approaches to more general supply chain structures.

9 Conclusions

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 supply chain formation protocol 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. We also identified Bayes-Nash equilibria policies for special classes of networks.

Further strategic analysis of the combinatorial auction is needed. Because the base strategic bidding policy we described is not generally a Bayes-Nash equilibrium, we would like to generalize or revise it to include Bayes-Nash equilibria for a broader class of networks. Even better would be a general method for automatically computing Bayes-Nash equilibria. But without strong guidance from the literature for how to approach the problem for general networks, finding, or even verifying Bayes-Nash equilibria is a challenging task. The quasi-equilibrium analysis we performed for one network may be a promising approach, but would be computationally expensive for larger networks. Alternatively, alternate auction types may reduce or eliminate the incentive to bid strategically in general supply chains, while maintaining reasonably high efficiency.

Using our present protocol, it would be straightforward to extend our supply chain model in various ways. Natural extensions could include consumer demand for multiple goods, alternative producer technologies, and multiple output choices, so long as the constraints can be expressed in linear terms. Introducing many choices would negatively affect the computational tractability of the auction. Moreover, with increased options in the network, agents would need to perform more Monte Carlo sampling, each requiring the solution of multiple combinatorial problems, to obtain comparable fidelity in computing bids.

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 his assistance in establishing the Bayes-Nash equilibrium results and David Parkes for constructive comments on an earlier version of this paper. 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.

References

[1] Arne Andersson, Mattias Tenhunen, and Fredrik Ygge. Integer programming for combinatorial auction winner determination. In *Fourth International Conference on Multi-Agent Systems*, pages 39–46, 2000.

[2] Moshe Babaioff and Noam Nisan. Concurrent auctions across the supply chain. In *Third ACM Conference on Electronic Commerce*, pages 1–10, 2001.

[3] Sushil Bikhchandani and John W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74:385–413, 1997.

[4] K. Chatterjee and W. Samuelson. Bargaining under incomplete information. *Operations Research*, 31:835–851, 1983.

[5] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, Fall 1971.

[6] William H. Davidow. *The Virtual Corporation: Structuring and Revitalizing the Corporation for the 21st Century*. HarperCollins Publishers, 1992.

[7] Sven de Vries and Rakesh Vohra. Combinatorial auctions: A survey. *INFORMS Journal on Computing*, to appear.

[8] Fcc 2001 combinatorial bidding conference. `http://wireless.fcc.gov/auctions/conferences/combin2001/index.html`.

[9] Drew Fudenberg and Jean Tirole. *Game Theory*. The MIT Press, 1998.

[10] Theodore Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.

[11] Frauk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95–124, 1999.

[12] Luke Hunsberger and Barbara J. Grosz. A combinatorial auction for collaborative planning. In *Fourth International Conference on MultiAgent Systems*, pages 151–158, 2000.

[13] David M. Kreps. *Game Theory and Economic Modelling*. Clarendon Press, Oxford, 1990.

[14] John Ledyard, Mark Olson, David Porter, and Joseph Sawnson. The first use of a combined value auction for transportation services. Technical Report 1093, Social Science Working Paper, March 2000.

[15] John Ledyard, David Porter, and Randii Wessen. A market-based mechanism for allocating space shuttle secondary payload priority. *Experimental Economics*, 2(3):173–195, March 2000.

[16] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. In *ACM Conference on Electronic Commerce*, pages 96–102, Denver, 1999.

[17] W. Leininger, P. Linhart, and R. Radner. Equilibria of the sealed-bid mechanism for bargaining with incomplete information. *Journal of Economic Theory*, 48:63–106, 1989.

[18] Jeffrey K. MacKie-Mason and Hal R. Varian. Generalized Vickrey auctions. Technical report, Dept. of Economics, Univ. of Michigan, July 1994.

[19] Thomas W. Malone and Robert J. Laubacher. The dawn of the e-lance economy. *Harvard Business Review*, pages 145–152, Sep-Oct 1998.

[20] R. Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56:434–450, 1992.

[21] R. Preston McAfee and John McMillan. Auctions and bidding. *Journal of Economic Literature*, 25:699–738, June 1987.

[22] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.

[23] David C. Parkes. *i*Bundle: An efficient ascending price bundle auction. In *ACM Conference on Electronic Commerce*, pages 148–157, 1999.

[24] David C. Parkes, Jayant Kalagnanam, and Marta Eso. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *Seventeenth International Joint Conference on Artificial Intelligence*, pages 1161–1168, 2001.

[25] David C. Parkes and Lyle H. Ungar. Iterative combinatorial auctions: Theory and practice. In *Seventeenth National Conference on Artificial Intelligence*, pages 74–81, 2000.

[26] David C. Parkes and Lyle H. Ungar. Strategic manipulation in iterative auctions: Proxy agents and price-adjustments. In *Seventeenth National Conference on Artificial Intelligence*, pages 82–89, 2000.

[27] 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.

[28] Michael H. Rothkopf, Aleksandar Pekeč, and Ronald M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44:1131–1147, 1998.

[29] Tuomas Sandholm and Fredrik 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.

[30] Mark A. Satterthwaite and Steven R. Williams. Bilateral trade with the sealed bid *k*-double auction: Existence and efficiency. *Journal of Economic Theory*, 48:107–133, 1989.

[31] Trade Extensions, Inc. Business case: Saving money for Volvo with a combinatorial auction. `http://www.tradeextensions.com`, 2001.

[32] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[33] William E. Walsh. *Market Protocols for Decentralized Supply Chain Formation*. PhD thesis, University of Michigan, 2001.

[34] William E. Walsh and Michael P. Wellman. A market protocol for decentralized task allocation. In *Third International Conference on Multi-Agent Systems*, pages 325–332, 1998.

[35] William E. Walsh and Michael P. Wellman. Efficiency and equilibrium in task allocation economies with hierarchical dependencies. In *Sixteenth International Joint Conference on Artificial Intelligence*, pages 520–526, 1999.

[36] William E. Walsh and Michael P. Wellman. MarketSAT: An extremely decentralized (but really slow) algorithm for propositional satisfiability. In *Seventeenth National Conference on Artificial Intelligence*, pages 303–309, 2000.

[37] William E. Walsh and Michael 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.

[38] William E. Walsh, Michael P. Wellman, and Fredrik Ygge. Combinatorial auctions for supply chain formation. In *Second ACM Conference on Electronic Commerce*, pages 260–269, 2000.

[39] Michael P. Wellman and Junling Hu. Conjectural equilibrium in multiagent learning. *Machine Learning*, 33:179–200, 1998.

[40] Michael P. Wellman, William E. Walsh, Peter R. Wurman, and Jeffrey K. MacKie-Mason. Auction protocols for decentralized scheduling. *Games and Economic Behavior*, 35(1/2):271–303, 2001.

[41] Peter R. Wurman. *Market Structure and Multidimensional Auction Design for Computational Economies*. PhD thesis, University of Michigan, 1999.

[42] Peter R. Wurman and Michael P. Wellman. A*k*BA: A progressive, anonymous-price combinatorial auction. In *Second ACM Conference on Electronic Commerce*, pages 21–29, 2000.

[43] Fredrik Ygge. *Market-Oriented Programming and its Application to Power Load Management*. PhD thesis, Lund University, 1998.