

Price Prediction Strategies for Market-Based Scheduling

Jeffrey K. MacKie-Mason Anna Osepayshvili Daniel M. Reeves Michael P. Wellman

University of Michigan

Ann Arbor, MI 48109 USA

{jmm, annaose, dreeves, wellman}@umich.edu

Abstract

In a market-based scheduling mechanism, the allocation of time-specific resources to tasks is governed by a competitive bidding process. Agents bidding for multiple, separately allocated time slots face the risk that they will succeed in obtaining only part of their requirement, incurring expenses for potentially worthless slots. We investigate the use of *price prediction* strategies to manage such risk. Given an uncertain price forecast, agents follow simple rules for choosing whether and on which time slots to bid. We find that employing price predictions can indeed improve performance over a straightforward baseline in some settings. Using an empirical game-theoretic methodology, we establish Nash equilibrium profiles for restricted strategy sets. This allows us to confirm the stability of price-predicting strategies, and measure overall efficiency. We further experiment with variant strategies to analyze the source of prediction's power, demonstrate the existence of self-confirming predictions, and compare the performance of alternative prediction methods.

Introduction

We address the design and evaluation of bidding strategies for participating in a market-based scheduling mechanism. Many resource-constrained scheduling problems are necessarily *decentralized*, due to the distribution of resources and information across autonomous and non-cooperative agents. One promising (and ubiquitous) approach is to allocate time-dependent resources through agent trading in market-based mechanisms. The success of an agent participating in a market-based scheduling mechanism hinges on implementing an effective trading strategy. Strategy choices also determine the success of the resulting schedules, and thus a strategic understanding of the problem informs the design and selection among alternative mechanisms.

In a decentralized problem, tasks and/or processors are associated with autonomous, independent entities (*agents*) who are self-interested (non-cooperative) and whose scheduling requirements or willingness to provide resources generate contention. The problem configuration is defined by the agents (each with resource endowments and preferences over possible allocations), and a resource allocation *mechanism*. The mechanism dictates the rules by which

agents interact (i.e., the set of messages that agents can exchange, typically via a *mediator*), and an allocation rule that maps the message space into allocations. An agent chooses a *strategy* that determines its messages in light of its objective function, and beliefs or knowledge about the endowments, preferences, and strategies of the other agents, as well as knowledge of the mechanism rules.

Most of the literature on market-based mechanisms for allocation problems addresses a particular problem: given the specification of the environment (agents with their endowments, preferences, information and beliefs, and the resources to be allocated, including any constraints), what is the mechanism that optimizes a social criterion subject to certain social constraints? To solve this problem, one typically assumes that agents will play their optimal (rational) strategy when participating in the mechanism.¹

Complex problems such as decentralized scheduling often call for multiple, dynamically interacting markets, rather than a single direct mechanism. For such complex mechanisms, except for the simplest of scheduling scenarios, it is generally not possible to solve analytically for optimal agent strategies. We therefore adopt a search-based approach, postulating strategy classes with well-motivated features or structures, and performing a computational game-theoretic analysis within the restricted classes identified.

The particular strategic feature we investigate in this paper is the use of economic price predictions to improve bidding strategies. The basic idea is that agents will typically have some knowledge about the demand and supply conditions for scheduled resources, from which they can derive estimates of the ultimate prices for these resources across time. The agents then use those predictions to guide their bidding behavior.

Our main research question is whether basing strategies on price predictions can improve bidder performance. When we find that it can, we further analyze the source of the improvement. We then explore several different price predictors to see whether our results are sensitive to the choice of prediction. Throughout we report the extent to which the new strategies lead to an increase or decrease in over-

¹This is the basic setup of the theory of optimal auctions (Myerson, 1981). Conitzer and Sandholm (2002) propose that one explicitly formulate and solve this optimization problem for particular allocation scenario instances.

all schedule efficiency. Our main finding is that a number of simple prediction methods, and simple strategy modifications to take advantage of the predictions, lead to a large increase in bidder performance at only a small cost in overall resource allocation efficiency.

Market-Based Scheduling Problem

In the simple scheduling problem we consider (Wellman et al., 2001), there are M units (called *time slots*) of a single schedulable resource, indexed $1, \dots, M$. Each of N agents has a single job that can be accomplished using the resource. Agent j 's job requires λ_j time slots to complete, and by accomplishing this job it obtains some value depending on the time it completes. Specifically, if j acquires λ_j time slots by deadline t , it accrues value $v_j(t)$. Deadline values are nonincreasing: $t < t'$ implies $v_j(t) \geq v_j(t')$.

In the *simultaneous ascending auction* for scheduling, a separate auction runs for each slot. Each auction can have multiple rounds of bidding. At any given time, the *bid price* on slot m is β_m , defined to be the highest bid b_j^m received thus far, or zero if there have been no bids. To be admissible, a bid must meet the bid price plus a bid increment (which we take to be one w.l.o.g.), $b_j^m \geq \beta_m + 1$. If an auction receives multiple admissible bids in a given round, it admits the highest (breaking ties arbitrarily). An auction is *quiescent* when a round passes with no new admissible bids. Each agent's payoff is then its deadline contingent valuation minus the total cost of all slots won.

The auctions proceed concurrently. When all are simultaneously quiescent, the auctions close and allocate their respective slots per the last admitted bids. Because no slot is committed until all are, an agent's bidding strategy on one slot cannot be contingent on the outcome for another slot. Thus, an agent j with $\lambda_j > 1$ inherently runs the risk—if it bids at all—that it will purchase some slots yet not acquire a sufficient set to accomplish its job. This is the well-known *exposure problem*, and arises whenever agents have positive interdependencies (*complementarities*) among goods allocated through separate markets. The exposure problem is perhaps the pivotal strategic issue in our market-based scheduling scenario, and motivates the approach we investigate here.

Price Prediction

The exposure problem presents a direct tradeoff. Bidding on a needed slot increases the prospects for completing a job, but also increases the loss in case the full set of required slots cannot be acquired. A rational bidding policy, therefore, would account for these expected costs and benefits, choosing to bid when the benefits prevail, and cutting losses in the alternative. Unfortunately, formulating and solving a complete model of the bidding decision does not appear to be tractable, due to the dynamic nature of the problem, and sensitivity to highly uncertain values and behaviors of other agents.² We therefore seek approximate strategies that

²In prior work (Reeves et al., to appear), we illustrate the subtle complexity of the problem, even when the other agents are known

capture some essential ingredients of the tradeoff, without incurring the computational burden of a full-fidelity analysis.

Consider the problem of agent 1, bidding against agent 2 in the situation described by Table 1. (Ignore agent 3 for the nonce.) Agent 2 needs only a single slot, and prefers earlier slots to later ones. It achieves a value of 8 if it completes its job with the first slot, 6 if the second, and 4 if the third. Agent 1 needs all three slots to complete its job, and achieves a value of 15 if successful. We assume the agents bid for these slots through a simultaneous ascending auction mechanism, with a bidding increment of 1.

Name	Job Length (λ)	$v(1)$	$v(2)$	$v(3)$
Agent 1	3	—	—	15
Agent 2	1	8	6	4
Agent 3	1	10	8	6

Table 1: A simple example illustrating the potential use of price prediction to address the exposure problem.

If the agents apply straightforward bidding strategies (defined below), one possible outcome (depending on random tie-breaking sequences) is that agent 2 wins the first slot at a price of 7, and agent 1 wins the second and third at prices of 5 and 3, respectively. This yields a surplus of 1 for agent 2, and -8 for agent 1. Agent 1 stops at this point under the straightforward bidding strategy, because there is no way it could complete its job profitably given the current state. However, it is clear from this example that dropping out is a mistake, as it could obtain slot 1 at a price of 8, for a net surplus of -1 , still a loss but a great improvement over the -8 it obtained under straightforward bidding.

It was precisely this sort of situation that led us in our earlier work (Reeves et al., to appear) to consider variations on straightforward bidding that took *sunk costs* into account. Intuitively, agent 1 should treat its commitment to buy slots 2 and 3 as a sunk cost, and continue to bid on slot 1 as long as that slot costs less than the value of its job. In this example, agent 1 will need to raise the bid just one more time, thus reducing its loss as described above. Moreover, the total value of jobs completed in this outcome is 15, compared to only 8 in the original solution.

However, treating the winning bids as sunk costs (even to a partial degree, as investigated in the previous paper) can be too indiscriminate. Consider the situation where agent 3 joins the bidding. With all three agents from Table 1 bidding straightforwardly, a possible outcome is that agent 3 wins the first slot at 7, agent 2 wins the second at 5, and agent 1 wins the third at 3. Again, agent 1 is caught by the exposure problem, stuck with a useless slot and a surplus of -3 . In this case, however, treating the cost as sunk is counterproductive. By continuing to bid on the first two slots it only increases its exposure, causing agent 3 to compete with it on the third slot as well. One possible outcome of this is that agent 1 ultimately wins all three slots, but at a total price of 23 (e.g., 10, 8, and 5—just sufficient to induce agent 3 to bid according to a straightforward policy.

drop out). This represents a surplus of -8 , worse than its baseline “sunk-unaware” outcome of -3 . The total value of jobs completed also decreases slightly, from 16 to 15.

Indeed, it may be too much to expect any bidding policy defined purely as a function from price quotes to bids to behave robustly as the mix of other agents changes. The effectiveness of a particular policy will in general be highly dependent on the characteristics of other agents in the environment. Unfortunately, a trading agent typically has at best gross distributional knowledge about the job lengths and deadline values of others. Even with such information, reasoning directly in terms of enumerations of bidding trajectories seems computationally infeasible. Thus, we seek summary constructs that can be applied directly to condition or guide relatively straightforward bidding behavior.

In particular, we consider information in the form of *price predictions*. In the scenario above, suppose agent 1 could predict before the auctions start that the prices would total 23 (for example, based on historical experience with similar situations). Then it could conclude that bidding is futile, not participate, and avoid the exposure problem altogether. Of course, we cannot expect in general to have perfect predictions. The hypothesis we investigate in this paper is that even relatively diffuse distributions can support predictions that significantly improve performance, even when employed in a straightforward manner.

Prior Work

Our scheduling model (single resource type, no job dependencies, no earliest-start constraints) is quite simple and well-understood, considered as a subclass of centralized scheduling problems (Tzafestas and Triantafyllakis, 1993; White, 1990). However, when the problem is augmented with decentralization constraints it is quite open (Wellman et al., 2001). To be sure, there is a substantial body of work on *distributed scheduling*—much of which considers far more elaborate scheduling models than that investigated here, including representative efforts from artificial intelligence (Sen and Durfee, 1998), distributed computing (Casavant and Kuhl, 1988), and operations management (Tharumarajah and Bemelman, 1997). Much of this research assumes (often quite reasonably given the setting of interest) a central algorithm designer, or that the agents share a common goal. Other work explores explicitly market-based mechanisms and the effects of particular bidding policies, but without scrutinizing the strategic problem from the agents’ perspective (Bertsekas, 1988; Walsh and Wellman, 2003).

We build on the economic literature about price-based allocation of combinatorial resources (Bikhchandani and Mamer, 1997; Bikhchandani and Ostroy, 2002; Gul and Stacchetti, 1999), and look to studies of simultaneous auction mechanisms (Milgrom, 2000; Peters and Severinov, 2001), for further insights about optimal (or equilibrium) trading strategies. However, our simple scheduling problem exhibits complementarities, in that for $\lambda > 1$, the value of individual slots is contingent on obtaining the others. This of course is the source of the exposure problem, discussed above, and the reason that straightforward bidding policies

do not suffice. Interestingly, auction theory (Krishna, 2002) to date has relatively little to say about how one *should* bid in simultaneous markets with complementarities, either in general or in the special case of problems with scheduling structure. This has led many researchers in recent years to explore *combinatorial auctions* (de Vries and Vohra, 2003), which allow bidders to express offers for bundles or combinations of goods explicitly, thus avoiding the exposure problem altogether.

Although we appreciate the virtues of combinatorial mechanisms, we are also aware of their limitations, most significantly their requirement for some competent authority to coordinate the allocation of interdependent resources, and the costs and delays associated with such coordination. It is a simple fact that today we see many markets operating separately, despite apparent strong complementarities for their respective goods. Whereas automation will very likely increase the prevalence of combinatorial markets, we expect that the issue of trading in separate yet dependent markets will remain.

Our methodology for exploring bidding strategies is based on our prior work (Reeves et al., to appear). Faced with an intractably large strategy space for a particular market game, we define a family of strategies parametrized by key strategic variables we can identify. We then employ sampling and simulation to construct an empirical normal-form game corresponding to expected payoffs for representative members of the restricted strategy class. Finally, we identify restricted equilibria by solving the generated game descriptions, exploiting symmetry or any other structure that may be available.

As noted above, our first study employing this methodology investigated the tradeoffs in treating current winnings as sunk costs. In the current work, we introduce the use of price prediction to guide bidding behavior. Price prediction played an important role in a recent market game involving complementary goods (Wellman et al., 2004), and of course price forecasting is generally known to be an important function in economic decision making.

Bidding Strategies

Straightforward Bidding

Our baseline for evaluating strategies is the *straightforward bidding* (SB) policy (Milgrom, 2000). A straightforward bidder bids myopically based on a best response to current price quotes. More precisely, the SB agent takes a vector of *perceived prices* (defined below) for the slots as given, and bids those prices for the bundle of slots that would maximize its surplus if it were to win all of its bids at those prices.

If agent j is assigned a set of slots X , it accrues value $v_j(X)$ based on the best deadline it can achieve: $v_j(X) = v_j(X(\lambda_j))$, where $X(t)$ is the t th time slot in X . Given that it obtains X at prices p , the agent’s *surplus* is its value less the amount paid, $\sigma(X, p) \equiv v_j(X) - \sum_{m \in X} p_m$. When agent j is winning the set of slots X_{-1} in the previous bidding round, we define the current *perceived prices* to be $\hat{p}_m = \beta_m$ for $m \in X_{-1}$, and $\hat{p}_m = \beta_m + 1$ otherwise. Then, under SB, agent j bids $b_j^m = \beta_m + 1$ for $m \in X^* \setminus X_{-1}$ such that

$$X^* = \arg \max_X \sigma(X, \hat{p}).$$

Straightforward bidders make no attempt to anticipate other agents' strategies. For instances with $\lambda_j = 1$, such anticipation is unnecessary, as agent j would not wish to change its bid even after observing what the other agents did (Bikhchandani and Mamer, 1997). If $\lambda_j = 1$ for all j , then SB leads to efficient allocations, up to an error proportional to the bid increment. Unfortunately, as illustrated above, the exposure problem destroys these positive results for cases with $\lambda_j > 1$ for some j .

Bidding with Price Prediction

As we have seen, straightforward bidding leads to neither optimal individual outcomes nor socially efficient allocations except in highly restricted problems. Our prior work (Reeves et al., to appear) explored one parametric generalization of SB, in which agents exhibit varying tendencies to treat provisional winning bids as sunk costs. In some settings, we indeed found that the "sunk-aware" modifications of SB produced some advantage. However, our simple example from Section "Price Prediction" demonstrates that sunk cost awareness is not a sufficient guide to optimal bidding.

One advantage of SB and the variants previously considered, is that they do not require any knowledge about the environment except for the current quotes. However, environmental factors are relevant, as indicated by the fact that different scheduling-problem configurations (numbers of agents and slots, job length and deadline-value distributions) often lead to qualitatively distinct equilibrium strategy profiles. Moreover, environmental information is often plausibly available, from knowledge of the structure of the market, or empirical observation of previous market outcomes (e.g., historical prices). In either case, the agent could (imperfectly) predict the final prices before the auctions start. We present examples of both analytic predictions from structural knowledge, and empirical predictions from observational data. However, our main emphasis is on how predictions might be used and evaluated, rather than how to derive them in the first place.

As noted above, when an agent's job length is one there is no exposure problem and straightforward bidding is a dominant strategy. Therefore the price prediction strategy reverts to straightforward bidding when $\lambda = 1$.

We assume that agents form their predictions before the auctions start and construct their bids based on these *predicted prices*, taking into consideration new information only if actual price quotes surpass the initial predictions. For example, if an agent's initial prediction on slot j 's final price was 8, but the current ask price for that slot in the auction has reached 9, the predicted price is effectively raised to 9.

We now formalize the price-predicting strategy family. Let $\pi = \pi_1, \dots, \pi_M$ be the vector of predicted prices. Define the perceived price vector, \hat{p} , as above for the SB strategy. We then define the *adjusted price prediction* for slot m , $\hat{\pi}_m$, as the maximum of predicted and perceived prices

$$\hat{\pi}_m \equiv \max(\pi_m, \hat{p}_m).$$

If $\lambda > 1$, the agent then chooses the set of slots on which to bid based on these adjusted predictions,

$$X^* = \arg \max_X \sigma(X, \hat{\pi}),$$

and issues bids for slots in X^* as in straightforward bidding.

The price-predicting strategy is thus parametrized by its predicted price vector, π . We label different predicted price vectors with a superscript, and then denote a specific price prediction strategy by $PP(\pi^x)$, where x labels particular predicted price vectors. Straightforward bidding is a special case of price prediction, with the predictions all equal to zero: $SB = PP(0)$. If the agent underestimates the final prices, it will behave identically to SB after the prices exceed the prediction. If the agent overestimates the final prices, it will stop bidding prematurely.

A Price Prediction Strategy Can Improve Performance

We first address our central question: Can agents using a price prediction strategy do better than agents using a benchmark strategy without price prediction? We find that minor strategy modifications based on unsophisticated price prediction information can improve performance substantially. We then measure the contribution of two separate aspects of the price prediction strategy to improved performance.

We conduct all of our analyses for a specific scheduling problem, employing a computational game-theoretic methodology developed in our prior work (Reeves et al., to appear).³ We assume there are five time slots available ($M = 5$), and that five agents are potentially interested in them ($N = 5$). Required job lengths are integers distributed uniformly on $[1, 5]$, and agent valuations for deadlines are integers distributed uniformly on $[0, 50]$. We adjust realized preference values to impose the constraint that later completion is not more valuable than earlier.

A Baseline Price-Prediction Strategy

We start by imagining that agents observe prices from prior auctions in which all players used the SB strategy, and use these average observed prices as their predictions. Let π^{BL} denote this baseline prediction (listed in Table 2).

To analyze the bidding policy based on this prediction, we consider a game restricted to the strategies SB and $PP(\pi^{BL})$. Through simulation, we generate the expected payoff matrix for this game; see Figure 1. In each column we record the average payoffs for each strategy in a particular profile. A profile is a list of strategies chosen by each of the five agents. The paired columns are given in lexicographic order from left to right, starting with all straightforward bidders (the benchmark profile): $\{SB, SB, SB, SB, SB\}$. The second column is $\{SB, SB, SB, SB, PP(\pi^{BL})\}$, and so forth. In the first profile, each SB agent receives an expected payoff of about 1.4. In the second profile, each SB agent receives an expected payoff of about 1.5, while the sole $PP(\pi^{BL})$ agent receives an expected payoff of about 2.3.

³The previous report includes many details of the methods, tangential for current purposes, that are omitted or glossed here.

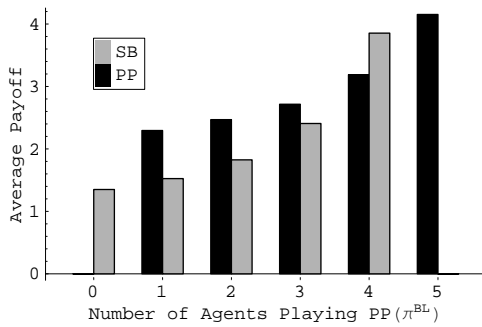


Figure 1: Payoff matrix when agents choose between two strategies: SB and PP(π^{BL}). The gray bars show the payoffs to SB agents and the solid bars show payoffs to predicting agents. The profiles are denoted by the number of predictors in the profile, from zero to five. Average payoffs were determined empirically from 200 thousand simulated games for each of six profiles.

We see from inspection of this payoff matrix that PP(π^{BL}) is a dominant strategy (since in every profile, the agents playing SB would do better switching to PP(π^{BL})) and thus also the unique symmetric equilibrium. The expected final prices are the average final prices in this equilibrium profile:⁴ (11.2, 6.8, 3.8, 2.0, 0.8). We summarize the results for this game (and those in subsequent sections) in Table 3. The efficiency relative to the optimal allocation is 86% (similarly computed, in general, from the efficiencies of each realized profile), which is 98% of the efficiency of the all-SB market. The average payoff in the equilibrium profile (again computed in general by appropriate mixing of the pure-strategy payoffs) is 4.2, which is 308% of the payoff to agents in the all-SB market. Thus, at a loss of only 2% in social efficiency, agents can improve their average performance by a factor of three if they use a simple price prediction based on average prices in an all-SB market.

Since all agents individually are better off in this equilibrium with price-predicting strategies, why is efficiency lower? The main gain to agents from using the PP(π^{BL}) price-predicting strategy is that they reduce the number of instances in which they are left paying for slots they cannot use (because they do not obtain a complete schedule). Although the exposure problem, as we showed above, can be very costly for *individual agents*, the allocation of these unused slots has no impact on social efficiency: if they are un-

⁴In general we compute this by weighting the final prices for each realized (pure strategy) profile by the probability of obtaining that profile given the equilibrium mixed-strategy probabilities. Given an N -player game with strategies S and all agents playing the mixed strategy $\{\alpha_1, \dots, \alpha_{|S|}\}$, the probability of a particular profile $(n_1, \dots, n_{|S|})$, where n_s is the number of players playing strategy s , is

$$\frac{N!}{n_1! \dots n_{|S|}!} \cdot \alpha_1^{n_1} \dots \alpha_{|S|}^{n_{|S|}}.$$

When there is a pure strategy equilibrium, we have the special case that all but one profile has zero weight.

used, it does not matter who “owns” the slot. The payment by the agent is just a transfer to the resource seller, and our calculation of efficiency is indifferent between whether the buyer or the seller has the slot or the money.⁵ Efficiency falls because to avoid the exposure problem, sometimes fewer schedules are completed, and so the allocation makes less valuable use of the available resources. In other words, price prediction prevents spurious purchases, improving buyers’ payoffs at substantial cost to the seller, netting a slight loss in social efficiency.

How Does Price Prediction Help?

We have shown that introducing price prediction can substantially improve expected payoffs for bidding agents. We now explore the reason for the improvement: under what conditions does our price-predicting agent bid differently than the benchmark agent, and how do these specific behavioral changes contribute to the improvement in expected payoffs?

Compared to benchmark bidding, our price-predicting agents use the prediction vector to modify two behaviors: choosing the best bundle of slots on which to bid, and deciding whether to participate in the bidding at all.⁶ We now decompose the effects reported in the previous subsection into those due to changes in participation and those due to changes in the choice of the bundle on which to bid.

To do this we construct a new bidding strategy. Agents first calculate the best bundle on which to bid using the same perceived prices as in the SB (no predictions) strategy. Then they choose to participate in the current round of bidding only if surplus from that bundle, valued at the *adjusted predicted prices*, is positive. Thus, the predicted prices are used only for the participation decision, and not for selecting the best bundle on which to bid.

In Figure 2 we present the payoff matrix for agents who choose either *participation-only* prediction, or straightforward bidding. The qualitative results for each possible strategy profile are similar to those in Figure 1. Again, the dominant (and therefore equilibrium) strategy is when all agents play PP(π^{BL}) with probability one. Expected payoff is 4.1, which is 2% lower than in the equilibrium with full-prediction agents as described above. Relative efficiency is slightly lower (half a percentage point).

To compare participation-only prediction to full prediction more thoroughly, we compute the ratio of payoff for participation-only prediction to payoff for full prediction for every possible environment of agents mixing between prediction and SB. We graph the results in Figure 3. The payoff to participation-only is less than to full-prediction as long as the probability of playing SB is less than 0.32. At best, if all

⁵In economics terminology, we are calculating a *Marshallian* money-metric welfare measure in which we equally weight the net surplus of all participants including buyers and sellers.

⁶We could say both behaviors manifest a single decision: on which bundle to bid, with “don’t participate” equivalent to bidding on the null bundle. We break the decision problem into participation and bundle selection because these two decisions are qualitatively different, and more closely match conventional understanding of auction behavior.

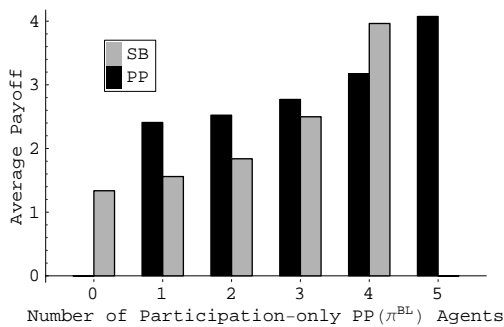


Figure 2: Payoff matrix when agents choose between SB and PP(π^{BL}) with participation-only prediction.




Figure 3: If agents play SB with probability less than 0.32, then full prediction does slightly better than participation-only prediction.

agents predict, participation-only prediction achieves 98.1% of the surplus of full prediction. At worst, in an environment of all SB agents, full prediction does slightly worse than participation-only prediction. Of course, we have found that in equilibrium agents will always use the prediction strategy.

Thus, we conclude that nearly all of the performance gain comes from the participation decision (that is, from sometimes dropping out earlier, which has the effect of reducing the risk of the exposure problem). In fact, far enough away from equilibrium the participation-only strategy secures more than 100% of the gains from price prediction.

Searching for Better Price Predictions

Any monotonically decreasing, nonnegative price vector is a candidate *prediction* vector. Agents can apply any reasonable functions to common information and obtain different predictions, or they might have access to different information. We have shown that one plausible but ad hoc price prediction vector can improve bidding performance substantially. We now address the further question: how well do strategies based on *different* prediction vectors perform?⁷

⁷We are interested in price predictions that are better only insofar as they lead to better performance for an agent bidding in a scheduling market. The quality of the prediction itself is not of any independent interest.




Figure 4: Convergence to self-confirming price-prediction vector, starting with initial prediction that all prices would be equal to zero. The prices at each iteration are determined by 500 thousand simulated games.

Self-Confirming Price Predictions

We first consider price predictions that are equal to the mean of the equilibrium prices if all agents use a price-predicting strategy. That is, if all agents use predictions, then the *self-confirming* predictions are those that on average *are correct*.

To solve for the self-confirming predictions we initialize the predicting agents with some prediction vector and simulate many games with the all-predict profile. When average prices obtained by these agents are determined, we replace the prediction vector with the average prices and repeat. When this process reaches a fixed point, we have the prices π such that when all agents play PP(π) the expected prices are π . We denote this prediction vector by π^{SC} . In Figure 4 we show the convergence to the five prices in π^{SC} . Within 30 iterations the prices have essentially converged, although there is some persistent oscillation. However, we found that by reseeding the prediction vector with the averages that the prices are oscillating around in Figure 4, this comprised a fixed point, which we used as π^{SC} .

We show the payoff matrix for the game with agents choosing between SB and PP(π^{SC}) in Figure 5. By inspection we see that PP(π^{SC}) is dominant: whatever strategies the first four agents play, the fifth agent always prefers PP(π^{SC}) to SB, and likewise for the other agents by symmetry.

The average final prices for the all-predict equilibrium are $\langle 13.1, 8.7, 5.4, 3.0, 1.2 \rangle$. The efficiency relative to the optimal allocation is 88%, which is negligibly higher than the efficiency of the all-SB market. The average payoff in the mixed-strategy equilibrium to an agent is 3.0, which is 226% of the payoff to agents in the all-SB market. Thus, at no loss in social efficiency, agents can more than double their average performance if they have accurate price predictions available.

Comparing Several Price Predictors

We consider two additional price prediction vectors. The baseline and self-confirming predictions were constructed

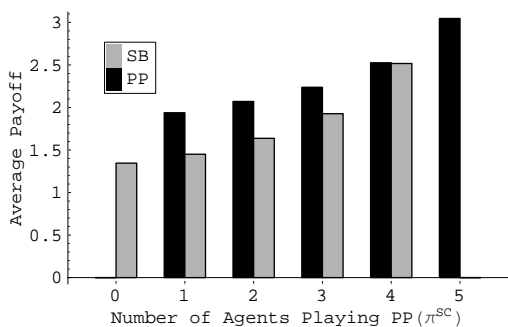


Figure 5: Payoff matrix for two strategies: straightforward bidding and price predicting with self-confirming predictions. Payoff matrix is constructed from 180 thousand games for each of the six profiles.

statistically, using observational data that might be available to agents. We now use another approach that we mentioned above: construct a model of the economic conditions and then solve analytically for equilibrium prices.

For our model, we suppose that the final prices form a competitive (or *Walrasian*) equilibrium in the scheduling market. This is guaranteed, for example, when straightforward bidders all demand only single slots, but is not true in general. However, in our experience, the final prices are generally not too far from competitive equilibrium prices. Therefore, we calculate the Walrasian equilibrium for our problem environment and use the resulting prices as predictions.

There is one problem we must overcome before calculating the Walrasian equilibrium: agents need to predict expected prices, which may deviate substantially from the Walrasian prices for any particular realization of preferences. We have been calculating expected payoffs by simulating a large number of market instances and averaging. We do likewise for expected prices. We calculate the *expected competitive equilibrium* prices, π^{ECE} by randomly generating a large number of game instances, analytically solving for the competitive equilibrium in each, and averaging across the resulting prices. We calculate the *expected demand competitive equilibrium* prices, π^{EDCE} , by calculating the expected demand function for each of many instances, and then solving for the competitive equilibrium based on the average demands. We label the corresponding bid strategies as $PP(\pi^{ECE})$ and $PP(\pi^{EDCE})$.

We now have five candidate price prediction strategies: SB (=PP(0)); $PP(\pi^{BL})$, $PP(\pi^{SC})$, $PP(\pi^{ECE})$, and $PP(\pi^{EDCE})$. The five π^x prediction vectors are given in Table 2. We consider a game in which each of the five agents can choose which of these five strategies to play. We present the resulting payoff matrix in Figure 6.

There are 126 possible combinations of five strategies among five players (each shown as a column in the payoff matrix). With this large problem, we did not find a dominant or asymmetric pure strategy Nash equilibrium through inspection. Instead, we used our method to solve for a




Figure 6: Payoff matrix for the game with four predicting strategies and SB. Each of the 126 columns (which are not meant to be distinguishable in this depiction) corresponds to a strategy profile: *all SB* ($\{SB, SB, SB, SB, SB\}$) through *all PP* (π^{BL}) in lexicographic order (given the ordering $\langle SB, PP(\pi^{EDCE}), PP(\pi^{ECE}), PP(\pi^{SC}), PP(\pi^{BL}) \rangle$). Payoffs were estimated by simulating 17 million games per profile.

Prediction Methods	Predicted Final Price Vectors				
SB (=PP(0))	0	0	0	0	0
$PP(\pi^{BL})$	14.8	10.7	7.6	4.6	1.9
$PP(\pi^{SC})$	13.0	8.7	5.4	3.0	1.2
$PP(\pi^{ECE})$	26.0	14.2	6.9	2.5	0.3
$PP(\pi^{EDCE})$	20.0	12.0	8.0	2.0	0.0

Table 2: Price predictions for all prediction methods. Compare to realized average prices in Table 3.

symmetric mixed strategy equilibrium. In the equilibrium we found, agents play $PP(\pi^{SC})$ with probability 0.45, and $PP(\pi^{EDCE})$ with probability 0.55.

The average equilibrium prices are $\langle 10.6, 6.5, 3.9, 2.2, 0.91 \rangle$. The expected payoff for an agent in the symmetric mixed strategy equilibrium is 4.3, which is 316% of the payoff in the all-SB market. The average efficiency of the allocation is 85%, which is 98% as efficient as the all-SB market. Thus, at a cost of only a 2% loss in social efficiency, agents can triple their average performance if they use a mixture of self-confirming and Walrasian predictions.

It may be surprising that all playing $PP(\pi^{SC})$ —and so predicting perfectly—is not an equilibrium. A price predictor with a perfect prediction has an incentive to make its prediction worse. The explanation is that the predictor’s performance is a function of how it uses the price prediction, as well as a function of its quality. The mixed equilibrium we found suggests that the accuracy of the self-confirming prediction is not sufficient for the strategy to prevail. How a prediction is incorporated into the strategy is as important, and exploring how one can make better use of a price prediction is an interesting topic for future research.

Games (i.e., strategy sets)	Equilibrium Profiles	% Eff.	Payoff	Average Final Price Vectors					
{SB, PP(π^{BL})}	all PP(π^{BL})	86	4.15	11.2	6.8	3.8	2.0	0.77	
{SB, PP(π^{BL}) w/ P.O.}	all PP(π^{BL}) w/ P.O.	85	4.07	11.8	6.9	3.7	1.7	0.58	
{SB, PP(π^{SC})}	all PP(π^{SC})	88	3.05	13.0	8.7	5.4	3.0	1.17	
{SB, PP(π^{BL}), PP(π^{SC}), PP(π^{ECE}), PP(π^{EDCE})}	0.45 SC, 0.55 EDCE	86	4.25	10.6	6.5	4.0	2.2	0.91	
Additional Profiles									
	all SB	87	1.35	14.8	10.7	7.6	4.6	1.90	
	all PP(π^{ECE})	74	5.80	4.7	2.1	1.7	1.2	0.55	
	all PP(π^{EDCE})	83	5.24	8.1	4.5	2.7	1.6	0.70	

Table 3: Average final price vectors, percent allocation efficiency relative to the global optimum, and average payoff for the (symmetric) equilibria of several games as well as for other, non-equilibrium profiles. PP(π^{BL}) w/ P.O. refers to participation-only prediction (see Section “How Does Price Prediction Help?”) using the baseline prediction vector. Sensitivity analysis (Reeves et al., to appear) confirms that all equilibrium results are robust to sampling noise and all reported efficiencies and payoffs have very small error bars.

Conclusion

We found that a simple modification of straightforward bidding, employing a prediction about final prices, can significantly improve bidder performance in a market-based scheduling environment. Predicted prices provide a guide to the agent about the exposure risk it incurs by bidding on the several slots it needs to complete a multi-period job. When it expects the relevant prices to exceed its job value, the agent refrains from bidding. At least for the environment we investigated, the binary choice of participation or not had a greater contribution to improvement than the use of predictions to guide choice of which slots to pursue.

The performance of a price prediction strategy depends, of course, on the quality of the prediction. We constructed five different predictions, and then solved a restricted game in which agents choose the prediction to use. One important prediction was self-confirming, in that when agents bid expecting these prices, they in fact result. Another important prediction was the Walrasian competitive equilibrium based on expected demand. In the mixed-strategy equilibrium of this game, agents combine with roughly equal probability the self-confirming prediction and the Walrasian prediction.

The performance of a price predictor also depends on how the predictions are used. The price-predicting strategies considered here are suboptimal. We consider the fact that even these simple PP(π^v) strategies offer such clearly superior performance over SB to be a strong confirmation of our main hypothesis. We expect that agents who predict distributions over prices, or otherwise more explicitly formulate the exposure tradeoff problem, can perform much better, and exploring this possibility is a subject for future work.

Although the results presented here are all specific to the particular scheduling problem and distribution of deadline-value preferences defined, we believe the general methods are broadly applicable. Both the SB and PP strategy families are well-defined for any simultaneous-auction environment. Whereas the exact price vectors derive from our particular problem, the price-prediction methods (BL, SC, ECE, EDCE) can likewise be applied to any given environment.

Finally, in presenting comparative strategy results, we do recognize that SB is to some extent a straw man, with no

advocates and well-known flaws. Unfortunately, advocates of *any* particular strategy for simultaneous interdependent auction problems are hard to find, as no bidding policy is known to perform generally well. Continued exploration of promising areas of the strategy space may help to remedy this situation.

Acknowledgments

This research has been supported in part by NSF grant CSS-9988715.

References

- D. Bertsekas. The auction algorithm: A distributed relaxation method for the assignment problem. *Annals of Operations Research*, 14:105–123, 1988.
- Sushil Bikhchandani and John W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74:385–413, 1997.
- Sushil Bikhchandani and Joseph M. Ostroy. The package assignment model. *Journal of Economic Theory*, 107:377–406, 2002.
- Thomas L. Casavant and Jon G. Kuhl. A taxonomy of scheduling in general-purpose distributed computing systems. *IEEE Transactions on Software Engineering*, SE-14(2):141–154, 1988.
- Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Eighteenth Conference on Uncertainty in Artificial Intelligence*, pages 103–110, Edmonton, 2002.
- Sven de Vries and Rakesh Vohra. Combinatorial auctions: A survey. *INFORMS Journal on Computing*, 15:284–309, 2003.
- Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95–124, 1999.
- Vijay Krishna. *Auction Theory*. Academic Press, 2002.

- Paul Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108:245–272, 2000.
- Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- Michael Peters and Sergei Severinov. Internet auctions with many traders. Technical report, University of Toronto, 2001.
- Daniel M. Reeves, Michael P. Wellman, Jeffrey K. MacKie-Mason, and Anna Osepayshvili. Exploring bidding strategies for market-based scheduling. *Decision Support Systems*, to appear.
- Sandip Sen and Edmund H. Durfee. A formal study of distributed meeting scheduling. *Group Decision and Negotiation*, 7:265–289, 1998.
- A. Tharumarajah and R. Bemelman. Approaches and issues in scheduling a distributed shop-floor environment. *Computers in Industry*, 34:95–109, 1997.
- S. Tzafestas and A. Triantafyllakis. Deterministic scheduling in computing and manufacturing systems: a survey of models and algorithms. *Mathematics and Computers in Simulation*, 35:397–434, 1993.
- William E. Walsh and Michael P. Wellman. Decentralized supply chain formation: A market protocol and competitive equilibrium analysis. *Journal of Artificial Intelligence Research*, 19:513–567, 2003.
- Michael P. Wellman, Daniel M. Reeves, Kevin M. Lochner, and Yevgeniy Vorobeychik. Price prediction in a trading agent competition. *Journal of Artificial Intelligence Research*, 21:19–36, 2004.
- Michael P. Wellman, William E. Walsh, Peter R. Wurman, and Jeffrey K. MacKie-Mason. Auction protocols for decentralized scheduling. *Games and Economic Behavior*, 35:271–303, 2001.
- K. P. White. Advances in the theory and practice of production scheduling. In C. T. Leondes, editor, *Advances in Control and Dynamic Systems*, pages 115–157. Academic Press, 1990.