Self-Confirming Price Prediction for Bidding in Simultaneous Ascending Auctions

Anna Osepayshvili, Michael P. Wellman, Daniel M. Reeves, and Jeffrey K. MacKie-Mason
University of Michigan
Ann Arbor, MI 48109 USA
{annaose, wellman, dreeves, jmm}@umich.edu

ABSTRACT

Simultaneous, separate ascending auctions are ubiquitous, even when agents have preferences over combinations of goods, from which arises the *exposure problem*. Little is known about strategies that perform well when the exposure problem is important. We present a new family of bidding strategies for this situation, in which agents form and utilize various amounts of information from predictions of the distribution of final prices.

The predictor strategies we define differ in their choice of method for generating the initial (pre-auction) prediction. We explore several methods, but focus on *self-confirming* predictions. An agent's prediction of characteristics of the distribution of closing prices is self-confirming if, when all agents follow the same predictor bidding strategy, the final price distributions that actually result are consistent with the utilized characteristics of the prediction.

We extensively analyze an auction environment with five goods, and five agents who each can choose from 53 different bidding strategies (resulting in over 4.2 million distinct strategy combinations). We find that the self-confirming distribution predictor is a highly stable, pure-strategy Nash equilibrium. We have been unable to find any other Nash strategies in this environment.

In limited experiments in other environments the self-confirming distribution predictor consistently performs well, but is not generally a pure-strategy Nash equilibrium.

1. SIMULTANEOUS ASCENDING AUCTIONS

A *simultaneous ascending auction* (SAA) (Cramton, 2005) allocates a set of M related goods among N agents via separate English auctions for each good. Each auction may undergo multiple rounds of bidding. At any given time, the *bid price* on good m is β_m, defined to be the highest non-repudiable bid b^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.), $b^m \geq \beta_m + 1$. 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.

The separate auctions proceed concurrently. When all are simultaneously quiescent, the auctions close and allocate their respective goods per the last admitted bids. Because no good is committed until all are, an agent's bidding strategy in one auction cannot be contingent on the outcome for another. Thus, an agent desiring a bundle of goods inherently runs the risk—if it bids at all—that it will purchase some but not all goods in the bundle. This is known as the *exposure problem*, and arises whenever agents have complementarities among goods allocated through separate markets. The exposure problem is the pivotal strategic issue in SAAs.

One approach to dealing with the exposure problem is to design mechanisms that take the complementarities directly into account. Much recent literature considers *combinatorial auctions* (Cramton et al., 2005; de Vries and Vohra, 2003), in which agents explicitly construct bids over bundles, and the auction mechanism determines optimal packages based on these bids. Although such mechanisms may provide an effective solution in many cases, there are often significant barriers to their application. Most significantly, combinatorial auctions require the existence of a competent authority to coordinate the allocation of interdependent resources, and incur 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, trading in separate dependent markets will remain fundamental.

Another approach to the exposure problem in an SAA is to design bidder strategies that perform well despite exposure risk. This is the approach we follow. The strategies we analyze share some common features. In each, an agent in a given auction round determines on which bundle to bid by calculating which would give the highest expected surplus according to some notion of expected prices or their distribution. Then the agent makes minimally incremental bids on the goods in the desired bundle that it is not already winning.

Formally, let $v_j(X)$ denote the value to agent j of obtaining the set of goods X. Given that it obtains X at prices $\boldsymbol{p}$, the agent's *surplus* is its value less the amount paid, $\sigma(X, \boldsymbol{p}) \equiv v_j(X) - \sum_{m \in X} p_m$. Then, each strategy we analyze is defined in terms of how the agent evaluates $\sigma(X, \boldsymbol{p})$ to select its currently preferred bundle X^*. Given that decision, in each strategy agent j bids $b_j^m = \beta_m + 1$ for the $m \in X^*$ that it is not already winning.

One widely-studied bidding strategy for SAAs is *straightforward bidding* (SB).[1] This strategy is quite simple, involving no anticipa-

[1] We adopt the terminology introduced by Milgrom (2000). The same strategy concept is also referred to as "myopic best response",

Dagstuhl Seminar Proceedings 05011
Computing and Markets
http://drops.dagstuhl.de/opus/volltexte/2005/202

tion of other agents' strategies. A straightforward bidder first calculates *perceived prices*. When agent j is winning the set of goods 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 selects the best bundle to maximize its surplus eavluated at the perceived prices, $X^* = \arg\max_X \sigma(X, \hat{\boldsymbol{p}})$, and then bids $b_j^m = \hat{p}_m$ for the goods $m \in X^*$ that it is not already winning.

Several results are known about the SB strategy. Agent j's value function exhibits *single-unit preference* iff for all X, $v_j(X) = \max_{m \in X} v_j(\{m\})$. If to the contrary, j's value for some bundle strictly exceeds that of its most valuable included singleton, we say the agent exhibits *bundle preference*. For an agent with single-unit preferences, SB is a *no regret* policy (Hart and Mas-Colell, 2000), as the agent would not wish to change its bid even after observing what the other agents did (Bikhchandani and Mamer, 1997).

When all agents have single-unit preference, and value every good equally, the situation is equivalent to a problem in which all buyers have an inelastic demand for a single unit of a homogeneous commodity. For this problem, Peters and Severinov (2001) showed that straightforward bidding is a perfect Bayesian equilibrium. Up to a discretization error, the allocations from SAAs are efficient when agents follow straightforward bidding. It can also be shown (Bertsekas, 1992; Wellman et al., 2001) that the final price vector will differ from the minimum unique equilibrium price by at most $\kappa \equiv \min(M, N)$. The value of the allocation, defined to be the sum of the bidder surpluses, will differ from the optimal by at most $\kappa(1+\kappa)$.

Unfortunately, the very nice properties for straightforward bidding with single-unit value do not carry over to bundle-preference problems. Indeed, the resulting price vector can differ from the minimum equilibrium price vector, and the allocation value can differ from the optimal, by arbitrarily large amounts (Wellman et al., 2001). More to the point, it is quite easy to construct examples where SB's obliviousness to the exposure problem causes the agent to incur significant losses in cases where these may have been anticipated and avoided. However, whereas the case against SB is quite clear, auction theory (Krishna, 2002) to date has relatively little to say about how one *should* bid in simultaneous markets with complementarities. In fact, determining an optimal strategy even when it is known that other agents are playing SB turns out to be an unsolved and surprisingly difficult problem, sensitive to the smallest details of preference distributions (Reeves et al., 2005).

Our gap in knowledge about SAA strategy is especially striking given the ubiquity of simultaneous auctions in economically significant settings. Indeed, markets for interdependent goods operating simultaneously and independently represents the normal or default state of affairs. Even for some markets that are expressly designed, most famously the US FCC spectrum auctions starting in the mid-1990s (McAfee and McMillan, 1996), a variant of the SAA is deliberately adopted, despite awareness of strategic complications (Milgrom, 2000). Simulation studies of scenarios based on the FCC auctions have shed light on some strategic issues (Csirik et al., 2001), as have accounts of some of the strategists involved (Cramton, 1995; Weber, 1997), but the general game is still too complex to admit definitive strategic recommendations.

2. BACKGROUND RESEARCH

Over the past several years, we have explored several SAA strategies in a market for scheduling resources (MacKie-Mason et al., 2004; Reeves et al., 2005). In the scheduling game, agents need to complete a job requiring a specified duration of resource, by acquiring the resource over individual time slots. The value for completing a job depends on when it is finished. Complementarities arise whenever jobs require more than a single time slot. The time slots are sold in simultaneous ascending auctions.

We have developed elsewhere an empirical methodology for exploring bidding strategies, which we describe and employ below (Section 5). In this approach, we develop a parametric description of a class of strategies, and explore the resulting strategy space through extensive simulation and analysis. In the prior work, we considered two extensions of SB designed to mitigate the exposure problem. First, we modified SB to approximately account for sunk costs, recognizing that goods an agent is already winning impose no incremental costs if other agents do not submit additional bids (Reeves et al., 2005). We implemented this strategy via a "sunk awareness" parameter ranging over [0,1], with zero treating all winning bids as sunk costs and one corresponding to unmodified SB. We then solved for settings of this parameter such that agents playing pure or mixed forms of this strategy are in Nash equilibrium. The equilibrium settings of this parameter are sensitive to the distribution of agent preferences (the value functions over subsets of goods), and we identified qualitatively distinct equilibria corresponding to different preference distributions.

For the second family of strategies we considered, agents select bundles based on predicted closing prices for each good (MacKie-Mason et al., 2004). We found that this approach is quite effective compared to the strategies based on a sunk-awareness parameter, including SB. Performance, of course, depends on the specific price prediction employed by the agent, as well as the distribution of agent preferences. Since prices from past transactions are observable, however, it is often plausible to estimate predictors directly from experience (real or simulated). Therefore, we defined strategies based on various *methods* for predicting, rather than based on particular numerical predictions. The set of prediction method strategies that survived in equilibrium in our experiments was relatively robust—compared to strategies based on a sunk-awareness parameter—to changes in the agent preferences. However, our analysis did not support conclusive statements about the general strategic stability of any particular strategy.

In this paper we build on our prior studies in several ways.

1. We extend the prediction-based strategies to employ *probability distributions* rather than simple point predictions.

2. We extend a notion of self-confirming prediction to the case of probability distributions, i.e., *self-confirming (SC) price distributions*, and present an iterative simulation-based algorithm to calculate these.

3. We conduct an extensive empirical game-theoretic analysis including the new SC distribution-based strategy, as well as a broad range of previously identified strategies. We extend our methodology to embrace large strategy sets without requiring exhaustive examination of the full combinatorial set of strategy profiles. Our key result is that playing the SC distribution-based strategy constitutes a pure-strategy, symmetric Nash equilibrium in the empirical game.

4. We examine the success of SC distribution-based bidding, and argue that—though not always optimal—it performs very well overall and is likely to be difficult to improve upon for general classes of SAAs.

or "myopically optimal", or even "myoptimal" (Kephart et al., 1998).

3. PROBABILISTIC PRICE PREDICTIONS

For agents with bundle preference, the exposure problem manifests in SAAs as a direct tradeoff. Bidding on a needed good increases the prospects for completing a bundle, but also increases the expected loss in case the full set of required goods 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.

3.1 Bidding with Price Prediction

Consider the $M = N = 3$ example presented in Table 1. Agents 2 and 3 have single-unit preferences. Agent 1 has bundle preference, and indeed needs all three goods to obtain any value.

Name	$v(\{1\})$	$v(\{2\})$	$v(\{3\})$	$v(\{1,2,3\})$
Agent 1	0	0	0	15
Agent 2	8	6	4	8
Agent 3	10	8	6	10

Table 1: An example of agent preferences.

If all three agents bid straightforwardly (SB), a possible outcome is that agent 3 wins the first good at 8, agent 2 wins the second at 5, and agent 1 wins the third at 3. Here, agent 1 is caught by the exposure problem, stuck with a useless good and a surplus of -3.[2] As we noted in Section 2, adding sunk-cost awareness to a bidding strategy can reduce exposure losses. However, as this example illustrates, not all exposure losses arise from improper handling of sunk costs. Suppose the first agent recognizes that its payment for good 3 is a sunk cost, and thus bids again on goods 1 and 2. If it could purchase all three at bids of $\{9, 6, 1\}$ its net loss would be reduced to -1. However, the other agents would continue bidding, and agent 1's ultimate loss would be greater than -3, so this "sunk aware" strategy would make it worse off.[3]

It may be too much to expect any bidding strategy defined purely as a function from price quotes to bids to behave robustly. The effectiveness of a particular strategy will in general be highly dependent on the characteristics of other agents in the environment. Often, however, a trading agent may have at least some beliefs about the distribution of other agent preferences. Thus, we turn to strategies that employ preference distribution beliefs to guide bidding behavior, rather than relying only on current price information as in the SB strategy.

One natural use for preference distribution beliefs is to form *price predictions* for the goods in an SAA. In the example above, suppose agent 1 could predict with certainty before the auctions start that the prices would total at least 16. Then it could conclude that bidding is futile, not participate, and avoid the exposure problem altogether. Of course, agents will not in general make perfect predictions. However, we find that even modestly informed predictions can significantly improve performance.

3.2 Point Price Prediction

We now define a strategy employing *point* estimates of final prices. The agent uses the point estimates to choose the bundle on which to bid in a given round, in contrast to SB which uses current perceived prices to select the preferred bundle for bidding. We first define the point estimates, and then how they are updated by this strategy when new information becomes available.

Let $\boldsymbol{\Omega}$ be the set of information available to an agent that is relevant to predicting the prices of the M goods. Partition $\boldsymbol{\Omega}$ as $(\boldsymbol{\Omega_0}; \mathrm{B})$, where B is information obtained during the course of the SAA (viz., the $t \times M$ history of bid prices, as of the tth round), and $\boldsymbol{\Omega_0}$ is information available to the agent prior to the auction. Let $\boldsymbol{\pi}(\boldsymbol{\Omega_0}, \mathrm{B})$ be a vector of predicted prices. Before the auction begins the price predictors are $\boldsymbol{\pi}(\boldsymbol{\Omega_0}, \phi)$, where ϕ is the empty set of auction bid information available pre-auction.

We next define how an agent following this strategy updates its predictions. In general, higher quotes should should cause agents to revise their predictions upwards. In particular, since the auction is ascending, once the current bid price for good m reaches β_m, there is zero probability that the final price $p_m < \beta_m$. Indeed, as a simplification, we define the strategy so that the only information used from the history of bid prices is that the predicted price can be no lower than the current bid price. In particular, to retain the myopia of the SB strategy, we define the current price prediction for good m to be the maximum of the initial prediction and the current perceived prices as defined for SB:[4]

$$\pi_m(\boldsymbol{\Omega_0}, \mathrm{B}) \equiv \max(\pi_m(\boldsymbol{\Omega_0}, \phi), \hat{p}_m).$$

Armed with these predictions, the agent chooses the set of goods on which to bid. If the agent has single-unit preference, it plays SB because that strategy is then optimal. Otherwise, similar to the SB strategy, the agent determines its preferred best bundle, but does so by evaluating bundles at the predicted prices:

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

where $\sigma(X, \boldsymbol{\pi})$ is the agent's surplus defined in Section 1. The agent then issues bids for goods in X^* as in straightforward bidding.

We thus have a family of point predicting strategies, parametrized by the vector of initial price predictions, $\boldsymbol{\pi}(\boldsymbol{\Omega_0}, \phi)$. We denote a specific point price prediction strategy by $PP(\boldsymbol{\pi}^x)$, where x labels particular initial prediction vectors.

3.3 Distribution-Based Price Prediction

Typically, the information available to an agent will support not just a point price prediction, but some estimate of additional characteristics of the stochastic distribution of final prices. Strategies using additional distribution information can at least weakly dominate strategies using only a predictor of the distribution mean.

This consideration motivates us to consider a family of strategies that employ more of an agent's beliefs about the distribution of prices. In this paper we specify a strategy that conditions on an agent's beliefs about the complete distribution of the final prices. We call this family *distribution-based price predictors*, or simply *distribution predictors*. We now define the distribution predictor, then explain how the initial distribution predictor is updated as new information is revealed through auction price quotes, and end by explaining how the agent uses the distribution predictions to select the bundle on which to bid in each round of the SAA.

Let $F \equiv F(\boldsymbol{\Omega_0}; \mathrm{B})$ denote an agent's belief about the joint cumulative distribution function over final prices. We assume that prices are bounded above by a known constant, V. Thus, the do-

[2] Depending on the sequence of bidding (when asynchronous), and the outcome of random tie-breaking (when synchronous) several different outcomes are possible, with agents following SB. All of them, however, leave agent 1 exposed, with negative surplus.

[3] For example, one possible outcome is that agent 1 wins all three goods at a total price of 23 (e.g., 10, 8, and 5—just sufficient to induce agent 3 to drop out). This represents a surplus of -8, worse than its baseline SB outcome of -3.

[4] We do not claim this updating process makes optimal use of the available information. Rather, it is simply the minimal adjustment consistent with the available observations.

main of F is $\{1, \ldots, V\}^M$. As before, we assume the agent generates $F(\mathbf{\Omega_0}; \phi)$, an initial, pre-auction predictor. We denote the strategy of bidding based on a particular distribution predictor by $PP(F^x)$, where x labels various distribution predictors.

We now turn to the prediction updating process. As with the point predictor, we restrict the updating in our distribution predictor to conditioning the distribution only on the fact that prices are bounded from below by $\boldsymbol{\beta}$.

Formally, let $\Pr(\boldsymbol{p}|\mathrm{B}; F)$ be the predicted probability, according to F, that the final price vector will be $\boldsymbol{p}$, conditioned on the information revealed by the auction B. Then, with $\Pr(\boldsymbol{p}|\phi; F)$ as the pre-auction initial prediction, we have:

$$\Pr(\boldsymbol{p}|\boldsymbol{\beta}; F) \equiv \begin{cases} \dfrac{\Pr(\boldsymbol{p}|\phi; F)}{\sum_{\boldsymbol{q} \geq \boldsymbol{\beta}} \Pr(\boldsymbol{q}|\phi; F)} & \text{if } \boldsymbol{p} \geq \boldsymbol{\beta} \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

(By $\boldsymbol{x} \geq \boldsymbol{y}$ we mean $x_i \geq y_i$ for all i.) Note that (1) is well defined for all possible $\boldsymbol{\beta}$ iff $\Pr(V, \ldots, V|\phi; F) > 0$.

We now describe how our distribution predictor strategy uses the prediction to choose bids. In general the agent would determine which bundle is expected to maximize surplus, with the expectation over final prices taken with respect to the joint distribution F:

$$X^* = \arg\max_X E_p\left[\sigma(X, \boldsymbol{p})|F\right].$$

However, calculating these expectations is very demanding. Therefore, for this paper we make a simplifying assumption. Our implementation of a distribution predictor strategy employs the predictions as if final prices are independent across goods. That is, we only use the information contained in the vector of marginal distributions, $(F_1, \ldots, F_M)$.

We follow a simple logic to exploit this reduced distribution information in the bid strategy. In SB the myopic agent calculates the best bundle evaluated at current perceived prices. Point predictors determine the best bundle evaluated at predictions of final prices. For the distribution predictor the agent determines the best bundle evaluated at *expected incremental* prices, taking the expectation with respect to the predicted price distribution, and calculating incremental prices by using the distribution to evaluate the degree to which the agent's current winning bids are likely to be sunk costs.

If an agent is currently not winning a good and bids on it, then it has no prior (sunk) commitment on that good, and if it wins, the incremental expected price is the full amount of the expected price. If the agent is currently winning a good, then the incremental expected price depends on the likelihood that the current bid price will be increased by another agent, so that the agent has to bid again to obtain the good; if it keeps the good at the current bid, the full price is sunk and thus does not affect the incremental cost of bidding.

For convenience, we define the expected price conditional on the most recent vector of bid prices, $\boldsymbol{\beta}$:

$$E(p_m|\boldsymbol{\beta}) = \sum_{q_m=0}^{V} \Pr(q_m|\boldsymbol{\beta}; F) q_m = \sum_{q_m=\beta_m}^{V} \Pr(q_m|\boldsymbol{\beta}; F) q_m,$$

where the first equality follows because the strategy treats prices as if they were independent.

We now define Δ_m, the expected incremental price for good m. If the agent is not winning good m, then the lowest final price at which it could is $\beta_m + 1$, and the expected incremental price is simply the expected price conditional on $\beta_m + 1$:

$$\begin{aligned} \Delta_m^{not-winning} &= E(p_m|\beta_m + 1) \\ &= \textstyle\sum_{q_m=\beta_m+1}^{V} \Pr(q_m|\beta_m + 1; F) q_m \end{aligned} \tag{2}$$

If the agent is winning good m, then the incremental price is zero if no one outbids the agent. With probability $1 - \Pr(\beta_m|\beta_m; F)$ the final price is higher than the current price, and the agent is outbid with a new bid price $\beta_m + 1$. Then, to obtain the good to complete a bundle, the agent will need to bid at least $\beta_m + 2$, and the expected incremental price is

$$\Delta_m^{winning} = (1 - \Pr(\beta_m|\beta_m; F)) \sum_{q_m=\beta_m+2}^{V} \Pr(q_m|\beta_m + 2; F) q_m \tag{3}$$

Combining (2) and (3) we get the complete expression for the expected incremental price for good m:

$$\Delta_m \equiv \begin{cases} (1 - \Pr(\beta_m|\beta_m; F)) \displaystyle\sum_{q_m=\beta_m+2}^{V} \Pr(q_m|\beta_m + 2; F) q_m & \text{if the agent is winning good } m \\ \displaystyle\sum_{q_m=\beta_m+1}^{V} \Pr(q_m|\beta_m + 1; F) q_m & \text{if the agent is not winning good } m \end{cases}$$

The final step in defining the strategy is to specify how an agent calculates its bids. To select the bundle on which it will bid, the agent evaluates bundle surplus at the expected incremental prices,

$$X^* = \arg\max_X \sigma(X, \boldsymbol{\Delta}).$$

The agent then issues bids for goods in X^* as in straightforward bidding.

4. SELF-CONFIRMING PRICE DISTRIBUTIONS

4.1 Definition and Existence

An *SAA environment* comprises an SAA mechanism over M goods, a set of N agents, and a probability distribution over M-good value functions for each agent. We now define a self-confirming price distribution for SAA environments in which all agents play the distribution predicting strategy we defined in the previous section.

DEFINITION 1. *Let SE be an SAA environment. The prediction F is a* self-confirming price distribution for SE *iff F is the distribution of prices resulting when all agents play bidding strategy $PP(F)$ in SE.*

We similarly refer to a prediction as *approximately self-confirming* if the definition above is satisfied for some approximate sense of equivalence between the outcome distribution and the prediction distribution.

The key feature of self-confirming prices, of course, is that agents make decisions based on predictions that turn out to be correct with respect to the underlying probability distribution.[5] Since agents are optimizing for these predictions, we might reasonably expect the strategy to perform well in an environment where its predictions are confirmed.

[5]In the economics literature an equilibrium with this feature is sometimes called a "fulfilled expectations equilibrium" (Novshek and Sonnenschein, 1982).

We also define a more relaxed version of self-confirmation. The actual joint distribution will in general have dependencies across prices for different goods. We are interested in the situation in which if the agents play a strategy based just on marginal distributions, that resulting distribution has the same marginals, despite dependencies.

DEFINITION 2. *Let SE be an SAA environment. The prediction $F = (F_1, \ldots, F_M)$ is a vector of* self-confirming marginal price distributions for SE *iff for all m, F_m is the marginal distribution of prices for good m resulting when all agents play bidding strategy $PP(F)$ in SE.*

Note that the confirmation of marginal price distributions is based on agents using these predictions as if the prices of goods were independent. However, we consider these predictions confirmed in the marginal sense as long as the results agree for each good separately, even if the joint outcomes do not validate the independence assumption.

A natural question to raise at this point is whether self-confirming predictions can actually be identified in plausible SAA environments. We demonstrate below that self-confirming marginal predictions can be found in some specific instances, at least approximately up to a numeric tolerance. However, it is easy to show that they cannot generally exist, by invoking a particular case known to be difficult for SAAs.

Specifically, consider the $M = N = 2$ configuration illustrated by Table 2. This is a version of the simple instance commonly employed to illustrate the absence of a competitive equilibrium (Cramton, 2005). There exist no prices for goods 1 and 2 such that both agents optimize their demands at the specified prices, and the markets clear.

Name	$v(\{1\})$	$v(\{2\})$	$v(\{1,2\})$
Agent 1	0	0	30
Agent 2	20	20	20

Table 2: A configuration with no price equilibrium.

PROPOSITION 1. *There exist SAA environments for which no self-confirming or marginally self-confirming price distributions exist.*

Proof. Define an SAA corresponding to the configuration of Table 2. Given a deterministic SAA mechanism,[6] for fixed value functions the outcome from playing any profile of deterministic trading strategies is a constant. Thus, the only possible self-confirming distributions (which were defined for agents playing the deterministic $PP(F)$ strategies) must assign probability one to the actual resulting prices. But given such a prediction, our trading strategy will pursue the agent's best bundle at those prices, and must actually get them since the prices are correct if the distribution is indeed self-confirming. But then the markets would all clear, contrary to the fact that the predicted prices cannot constitute an equilibrium, since such prices do not exist in this instance. □

Despite this negative finding, we conjecture that for a large class of nondegenerate preference distributions, self-confirming price distributions or close approximations thereof will actually exist, and can be computed given a specification of the preference distribution. We present a procedure for deriving such distributions, and some evidence for its effectiveness, in the sections below.

4.2 Deriving Self-Confirming Price Distributions

Given an SAA environment—including the distributions over agent preferences—we derive self-confirming price distributions through an iterative simulation process. Specifically, we start from an arbitrary prediction F^0, and run many instances of the SAA environment (sampling from the given preference distributions) with all agents playing strategy $PP(F^0)$. We record the resulting prices from each instance, and designate the sample distribution observed by F^1. We then run another battery of instances, or *iteration*, with agents playing $PP(F^1)$, and repeat the process in this manner for some further series of iterations. If it ever reaches an approximate fixed point, with $F^t \approx F^{t+1}$ for some t, then we have statistically identified an approximate self-confirming price distribution for this environment. (Due to sampling error, the approximate version of the concept is the best we can attain through simulation.)

Any reasonable measure of similarity of probability distributions combined with a threshold constitutes an operable policy for validating approximate self-confirmation. We employ the Kolmogorov-Smirnov (KS) statistic, defined as the maximal distance between any two corresponding points in the CDFs:

$$KS(F, F') = \max_x |F(x) - F'(x)|.$$

When we seek self-confirmation only of predictors for the marginal distributions, we measure KS distance separately for each good, and take the largest value. That is, we define $KS_{\text{marg}} = \max_m KS(F_m, F'_m)$.

Our complete procedure for deriving approximate self-confirming price distributions is defined by specifying:

1. a number of samples per iteration,
2. a threshold on KS or KS_{marg} on which to halt the iterations and return a result,
3. a maximum number of iterations in case the threshold is not met,
4. a smoothing parameter k designating a number of iterations to average over when the procedure reaches the maximum iterations without finding an approximate fixed point.

The bound on the number of iterations ensures that this procedure terminates and returns a price distribution, which may or may not be self-confirming. When this occurs, the smoothing parameter avoids returning a distribution that is known to cause oscillation. However, the apparent nonexistence of a self-confirming equilibrium in this case suggests the problem cannot be totally eradicated, and we do not expect the strategy to perform as well when the underlying oscillations are large.

To illustrate the process, we present the calculation for an example. First we specify an SAA environment that represents a market-based scheduling problem. There are five agents competing for five time slots. An agent's value function is defined by its job length and its value for meeting various deadlines. We drew job lengths randomly from $U[1, 5]$. We chose deadline values randomly from $U[1, 50]$ then pruned to impose monotonicity; for details see Reeves et al. (2005).

[6] That is, without any nondeterministic effects of asynchrony or random tie-breaking. For any reasonable degree to which these are present, we can construct a more extreme example that would swamp any noise from these effects.

Good	Mean Price	Standard Deviation
1	10.8	7.7
2	6.5	5.1
3	4.1	3.7
4	2.3	2.5
5	1.0	1.3

Sample size per iteration: one million.

Table 3: Descriptive statistics for a self-confirming price distribution calculated in six iterations.

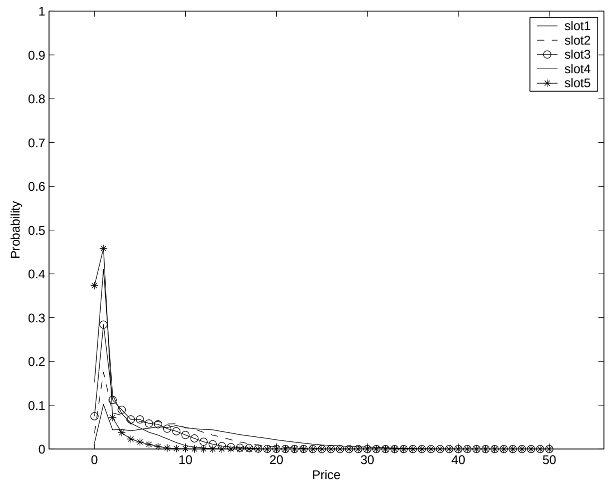

Figure 1: Self-confirmed probability distribution functions for five goods, calculated in five iterations.

Next we set the parameters for for our algorithm: one million games per iteration, and a 0.01 KS convergence criterion.[7] We ran the algorithm, playing the million games per iteration to generate an empirical price distribution.

The predicted and empirical distributions quickly converged, with a KS statistic of 0.007 after only six iterations. We report descriptive statistics for the result in Table 3, and we plot the SC price distributions for the five slots in Figure 1.

4.3 Experiments in Iterative SC Price Estimation

We described above one example of an SC price distribution. To test the hypothesis that iterative derivation will produce useful results with some regularity, we applied the procedure to the several variations on this SAA environment.

Specifically, we defined 23 instances of the market-based scheduling problem, for varying numbers of agents and goods. We again drew deadline values from $U[1, 50]$ and pruned them for monotonicity. We used two probability models for job lengths in the first 19 instances. In the *uniform* model, they are drawn from $U[1, M]$. In the *exponential* model job length λ has probability $2^{-\lambda}$, for $\lambda = 1, \ldots, M - 1$, and probability $2^{-(M-1)}$ when $\lambda = M$.

We applied our iterative SC derivation algorithm with the same parameters as above to 11 instances of the uniform model, comprising some combination of $3 \leq N \leq 9$ and $3 \leq M \leq 7$. In each case, the procedure found a set of approximately self-confirming marginal price distributions (KS threshold 0.01) within 11 iterations. Similarly, for 11 instances of the exponential model, with the number of agents and goods varying over the same range, we found SC distributions within seven iterations. Thus, for this class of environments the iterative method is quite successful.

The 23d instance was designed to be more challenging: the $N = M = 2$ example with fixed preferences described Table 2. Since there exists no SC distribution, our algorithm did not find one, and as expected after a small number of iterations it began to oscillate among a few states indefinitely. After reaching the limit of 100 iterations, our algorithm returned as its prediction distribution the average over the last $k = 10$.

[7]Since KS is a distance between CDFs, a 0.01 threshold is equivalent to a maximal one percentage point probability difference at any point in the two distributions.

5. EMPIRICAL GAME ANALYSIS

We now analyze the performance of self-confirming distribution predictors in a variety of SAA environments, against a variety of other strategies. We first describe our empirical game analysis method and the subset of strategy space we explore. Then we present the results.

5.1 Method

We briefly describe a computational methodology for analyzing games and strategies, extending the approach developed in our prior work (Reeves et al., 2005; MacKie-Mason et al., 2004), and other recent studies in a similar empirical vein (Armantier et al., 2000; Kephart et al., 1998; Walsh et al., 2002). See MacKie-Mason and Wellman (2005) for a more complete description.

We use an empirical methodology for environments that are not solvable using analytic methods. In general, no analytic solutions have been found for SAAs with complementary goods. Exhaustively analyzing possible strategies by enumeration is computationally intractable: for an SAA with incomplete information the strategy space is the set of mappings from all price quote histories and own preferences (itself potentially as large as the space of functions from the powerset of goods to the reals) to bid vectors. Our empirical method applies computational methods (due to analytic intractability) to a restricted strategy space (due to computational intractability) to find equilibrium strategies and to analyze their characteristics.

Our methodology invokes the following steps:

1. **Specify an environment.**
2. **Generate candidate strategies.**
3. **Estimate the "empirical game".**
4. **Analyze the empirical game.**

The remaining sections describe how we implemented these steps to evaluate the performance of the self-confirming (SC) distribution prediction strategy, $PP(F^{SC})$.

5.2 Environments considered

We studied SAAs applied to market-based scheduling problems. We defined examples of such problems in Sections 4.2 and 4.3. Particular environments are defined by specifying the number M of goods, the number N of agents, and a preference distribution which for the scheduling problem comprises a probability distribution over job lengths and another distribution over deadline values. We provide details below where we present the analyses.

The bulk of our computational effort went into an extensive analysis of one particular environment, the $N = M = 5$ uniform model presented above. As described in Section 5.5, the empirical game for this setting provides much evidence supporting the unique strategic stability of $PP(F^{SC})$. We complement this most detailed trial with smaller empirical games for a range of other scheduling-based SAA environments. Altogether, we have studied selected environments with uniform, exponential, and fixed distributions for job lengths; a modified uniform distribution for deadline values; and agents in $3 \leq N \leq 9$; goods in $3 \leq M \leq 7$.

5.3 Strategy space explored

To varying degrees, we have analyzed the interacting performance of 53 different strategies. These were drawn from the three strategy families described above: SB, point predictor, and distribution predictor. For each family we varied a defining parameter to generate the different specific strategies. In summary, we considered:

- SB and 20 SB variants, by varying the degree of sunk awareness (Reeves et al., 2005) by increments of 0.05 on $[0, 1]$.
- 13 point predictor strategies (MacKie-Mason et al., 2004), by varying the method used to generate the prediction vector, $\pi_m(\mathbf{\Omega_0}, \phi)$.
- 19 distribution predictor strategies, by varying the method used to generate the distribution prediction, $F(\mathbf{\Omega_0}; \phi)$.

Naturally, our emphasis is on evaluating the performance of $PP(F^{SC})$ in combination with the other strategies. One of the noteworthy alternatives is $PP(F^{SB})$, which employs the price distribution prediction formed by estimating (through simulation) the prices resulting from all agents playing SB.

An exhaustive description of the prediction methods used for point and distribution predictors is too long for this paper.[8] To summarize, we obtained predictions from variations on several basic methods: simulating an environment with all agents playing SB; by solving for the competitive equilibrium of the SAA game; and solving for self-confirming predictors (see Section 4.2). For distribution predictors we also generated degenerate and Gaussian distributions based on point predictor vectors.

5.4 Estimate and solve the empirical game

The next step in our empirical method is to convert the SAA in extensive form, which is not tractable, to a manageable game in normal form. We use Monte Carlo simulation methods to estimate the *payoff function* mapping profiles of agent strategies to expected payoffs for each agent. That is, given a profile of strategies followed by other agents, we repeatedly draw preferences and assign them to agents, simulate the auction protocol for the given strategy profile to quiescence, and average the resulting surpluses to estimate the expected payoffs to that strategy profile. We continue this Monte Carlo procedure for as many strategy profiles as we care to analyze.

Our environments are symmetric in the strategies available to agents, and in the payoffs received by agents, so the payoff function is also symmetric. Given N agents and S possible strategies, the number of distinct strategy profiles is $\binom{N+S-1}{N}$. For our primary example below there are five agents, thus there are over four million different strategy profiles to evaluate. Since we determine the expected payoffs empirically for each profile by running millions of simulations of the auction protocol, estimating the entire payoff function is infeasible. However, we can estimate the complete payoff matrix for various subsets of our 53 strategies. And as we describe below, we do not need the full payoff matrix to reach conclusions about equilibria in the 53-strategy game.

Given a complete payoff function over some subset of strategies, we solve for pure and mixed strategy equilibria using a variety of techniques, including replicator dynamics, function minimization, and the algorithms in the GAMBIT (McKelvey et al., 1992) game-solver library.

5.5 5×5 Uniform Environment

By far the largest empirical SAA game we have constructed is for the SAA scheduling environment discussed in Section 4.2, with five agents, five goods, and uniform distributions over job lengths and deadline values. We have estimated payoffs for 4916 strategy profiles, out of out of the 4.2 million distinct combinations of 53 strategies. Payoff estimates are based on an average of 10 million samples per profile (though some profiles were simulated for as few as 200 thousand games, and some for as many as 200 million). Despite the sparseness of the estimated payoff function (covering only 0.1% of possible profiles), we have been able to obtain several results.

First, as discussed above, we conjectured that the self-confirming distribution predictor strategy, $PP(F^{SC})$, would perform well. We have directly verified this: *the profile where all five agents play a pure $PP(F^{SC})$ strategy is a Nash equilibrium*. That is, we verified that no unilateral deviation to any of the other 52 pure strategies is profitable. Note that in order to verify a pure-strategy symmetric equilibrium (all agents playing s) for N players and S strategies, one needs only S profiles: one for each strategy playing with $N-1$ copies of s. Similarly, to refute the possibility of a particular profile being in Nash equilibrium, we need find only one profitable deviation profile (i.e., obtained by changing the strategy of one player to obtain a higher payoff given the others).

The fact that $PP(F^{SC})$ is pure symmetric Nash for this game does not of course rule out the existence of other Nash equilibria. Indeed, without evaluating any particular profile, we cannot eliminate the possibility that it represents a (non-symmetric) pure-strategy equilibrium itself. However, the profiles we did estimate provide significant additional evidence, including the elimination of broad classes of potential symmetric mixed equilibria.

Let us define a strategy *clique* as a set of strategies for which we have estimated payoffs for all combinations. Each clique defines a subgame, for which we have complete payoff information. Within our 4916 profiles we have eight maximal cliques, all of which include strategy $PP(F^{SC})$.[9] For each of these subgames, we have used the Gambit algorithms to determine that $PP(F^{SC})$ is the only strategy that survives iterated elimination of (strictly) dominated (pure) strategies (IEDS).[10] It follows that $PP(F^{SC})$ is the unique (pure or mixed strategy) Nash equilibrium in each of these clique games.

From this result, we can further conclude that in the full 53-strategy game there are no equilibria with support contained within any of the cliques, other than the special case of the pure-strategy $PP(F^{SC})$ equilibrium.[11]

[8] An appendix is in preparation. Details are available from the authors on request.

[9] One clique is a 10-strategy game with 2002 unique profiles; three are 5-strategy games (126 profiles each); one is a 4-strategy game (56 profiles); and three are 3-strategy games (21 profiles).

[10] Even better, $PP(F^{SC})$ is a dominant strategy in three of the subgames (two 3-strategy subgames and the one 5-strategy subgame in which it appears).

[11] By the IEDS result, any agent that plays a mixture of any clique strategy set can do better by deviating and playing $PP(F^{SC})$.

Analysis of the available two-strategy cliques (not generally maximal) provides further evidence about potential alternative equilibria. Of the $\binom{52}{2} = 1326$ pairs of strategies not including $PP(F^{SC})$, we have all profile combinations for 46. Based on profiles estimated, we can derive a lower bound of 0.32 on the value of ϵ such that a mixture of one of these pairs constitutes an ϵ-Nash equilibrium. In other words, we have determined that for any symmetric profile defined by such a mixture, an agent can improve its payoff by a minimum of 0.32 through deviating to some other pure strategy. For reference, the payoff for the all-$PP(F^{SC})$ profile is 4.51, so this represents a nontrivial difference.

Finally, for each of the 4916 evaluated profiles, we can derive a bound on the ϵ rendering the profile itself an ϵ-Nash pure-strategy equilibrium. The three most strategically stable profiles by this measure are:

1. all $PP(F^{SC})$: $\epsilon = 0$ (confirmed Nash equilibrium)
2. one $PP(F^{SB})$, four $PP(F^{SC})$: $\epsilon > 0.13$
3. two $PP(F^{SB})$, three $PP(F^{SC})$: $\epsilon > 0.19$

All the remaining profiles have $\epsilon > 0.25$ based on confirmed deviations.

Our conclusion from these observations is that $PP(F^{SC})$ is a highly stable strategy within this strategic environment, and likely uniquely so. Of course, only limited inference can be drawn from even an extensive analyis of only one particular distribution of preferences.

5.6 Self-confirming prediction in other environments

In order to test whether the strong performance of $PP(F^{SC})$ generalizes across other SAA environments, we performed smaller versions of the empirical game analysis on variations of the model above. Specifically, we explored 13 instances of the market-based scheduling problem:

- nine with the uniform model, with 3–8 agents and 3–7 goods
- three with the exponential model, with 3–8 agents and 3–5 goods
- one with fixed preferences, corresponding to the counterexample model of Table 2

For each we derived self-confirming price distributions (failing in the last case, of course), as reported in Section 4.3. We also derived price vectors and distributions for the other prediction-based strategies. For each symmetric game (the uniform and exponential models), we evaluated 27 profiles: one with all $PP(F^{SC})$, and for each of 26 other strategies s, one with $N-1$ $PP(F^{SC})$ and one s. For the non-symmetric game with fixed preferences, we evaluated all 53 profiles with at least one agent playing $PP(F^{SC})$. We ran between one and seven million games per profile in all of these environments.

In ten out of the twelve symmetric and uncertain environments, $PP(F^{SC})$ was verified to be an ϵ-Nash equilibrium for $\epsilon < 0.1$. In only two of these (one each uniform and exponential), however, was it an exact equilibrium. The two worst environments of the twelve were uniform with $N = 3$ and $M = 5$ $\epsilon = 0.10$, and uniform with $N = 8$ and $M = 7$ $\epsilon = 0.14$. In the last case, expected payoff for all-$PP(F^{SC})$ was 2.67, so ϵ represents about 5% of the value. For no other case did it reach 2%. Overall, we regard this as favorable evidence for the $PP(F^{SC})$ strategy across the range of market-based scheduling environments, though clearly not as strong as our results from the large-scale analysis of the uniform 5×5 case. Further exploration of these cases (i.e., evaluating more profiles) may shed light on the true strategic relation of $PP(F^{SC})$ to the other strategies where it is not apparently a pure-strategy Nash equilibrium.

Not surprisingly, the environment with fixed preferences is an entirely different story. Recall that in this case the iterative procedure failed to find a self-confirming price distribution. The distribution it settled on was quite inaccurate, and the trading strategy based on this performed poorly—generally obtaining negative payoffs regardless of other strategies. Since one of the available strategies simply does not trade, $PP(F^{SC})$ is clearly not a best-response player in this environment.

6. DISCUSSION

We have presented a general trading strategy for SAA environments that places bids based on probabilistic predictions of final good prices. Such a policy tackles the exposure problem head-on, by explicitly weighing the risks and benefits of placing bids on alternative bundles, or no bundle at all. This strategy generalizes our previous work on bidding based on point price predictions, and like that scheme is parametrized by the *method* for generating predictions. Given methods that take contextual conditions into account (e.g., distributions of agent preferences), particular trading strategies are potentially robust across varieties of SAA environments.

The method we consider most promising employs what we call *self-confirming price distributions*. A price distribution is self-confirming if it reflects the prices generated when all agents play the trading strategy based on this distribution. Although such self-confirming distributions may not always exist, we expect they will (at least approximately) in many environments of interest, especially those characterized by relatively diffuse uncertainty and a moderate number of agents. An iterative simulation algorithm appears effective for deriving such distributions.

Given the analytic and computational intractability of the game induced by an SAA environment, we evaluated our approach using an empirical game-theoretic methodology. We explored a restricted strategy space including a range of candidate strategies identified in prior work. Despite the infeasibility of exhaustively exploring the profile spaces, our analyses support several game-theoretic conclusions. The results provide favorable evidence for our method—very strong evidence in one environment we investigated intensely, and somewhat weaker (in some respects mixed) evidence for a range of variant environments.

The self-confirming price-prediction trading strategy shares features with other bidding policies proposed in the literature. Greenwald and Boyan (2004) explicitly employ distributions of prices, and consider approaches to optimize choice of bundles given general joint price distributions. They evaluate their method on simultaneous one-shot auctions, as well as in an SAA-like environment included in the TAC travel shopping game (Wellman et al., 2003). The strategy proposed by Gjerstad and Dickhaut (1998) for trading in continuous double auctions employs price distributions for determining what to bid, and calibrates these distributions online based on experience in the dynamic auction.

Neither the strategy we present here nor any other strategy is likely to be universally best across SAA environments. Nevertheless, we conjecture that the self-confirming price prediction strategy will be difficult to beat by much for any broad scenario class. If agents make optimal decisions with respect to prices that turn out to be right, there may not be room for performing a lot better. On the other hand, there are certainly areas where improvement should be possible, for example:

- accounting for one's own effect on prices
- incorporating price dependencies with reasonable computational effort
- more graceful handling of instances without self-confirming price distributions
- timing of bids: trading off the risk of premature quiescence with the cost of pushing prices up

We intend to explore some of these opportunities in follow-on work in this domain.

Acknowledgments

Rahul Suri contributed to the equilibrium analysis in Section 5.

References

Olivier Armantier, Jean-Pierre Florens, and Jean-Francois Richard. Empirical game-theoretic models: Constrained equilibrium and simulations. Technical report, State University of New York at Stonybrook, 2000.

Dimitri P. Bertsekas. Auction algorithms for network flow problems: A tutorial introduction. *Computational Optimization and Applications*, 1:7–66, 1992.

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

Peter Cramton. Money out of thin air: The nationwide narrowband PCS auction. *Journal of Economics and Management Strategy*, 4:267–343, 1995.

Peter Cramton. Simultaneous ascending auctions. In Cramton et al. (2005).

Peter Cramton, Yoav Shoham, and Richard Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2005.

János A. Csirik, Michael L. Littman, Satinder Singh, and Peter Stone. FAucS: An FCC spectrum auction simulator for autonomous bidding agents. In *Second International Workshop on Electronic Commerce*, volume 2232 of *Lecture Notes in Computer Science*, pages 139–151. Springer-Verlag, 2001.

Sven de Vries and Rakesh Vohra. Combinatorial auctions: A survey. *INFORMS Journal on Computing*, 15(3):284–309, 2003.

Steven Gjerstad and John Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.

Amy Greenwald and Justin Boyan. Bidding under uncertainty: Theory and experiments. In *Twentieth Conference on Uncertainty in Artificial Intelligence*, pages 209–216, Banff, 2004.

Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68:1127–1150, 2000.

Jeffrey O. Kephart, James E. Hanson, and Jakka Sairamesh. Price and niche wars in a free-market economy of software agents. *Artificial Life*, 4:1–23, 1998.

Vijay Krishna. *Auction Theory*. Academic Press, 2002.

Jeffrey K. MacKie-Mason, Anna Osepayshvili, Daniel M. Reeves, and Michael P. Wellman. Price prediction strategies for market-based scheduling. In *Fourteenth International Conference on Automated Planning and Scheduling*, pages 244–252, Whistler, BC, 2004.

Jeffrey K. MacKie-Mason and Michael P. Wellman. Automated markets and trading agents. In *Handbook of Agent-Based Computational Economics*. Elsevier, 2005.

R. Preston McAfee and John McMillan. Analyzing the airwaves auction. *Journal of Economic Perspectives*, 10(1):159–175, 1996.

Richard D. McKelvey, Andrew McLennan, and Theodore Turocy. Gambit game theory analysis software and tools, 1992. `http://econweb.tamu.edu/gambit`.

Paul Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108:245–272, 2000.

William Novshek and Hugo Sonnenschein. Fulfilled expectations Cournot duopoly with information acquisition and release. *Bell Journal of Economics*, 13:214–218, 1982.

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*, 39:67–85, 2005.

William E. Walsh, Rajarshi Das, Gerald Tesauro, and Jeffrey O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, Edmonton, 2002.

Robert J. Weber. Making more from less: Strategic demand reduction in the FCC spectrum auctions. *Journal of Economics and Management Strategy*, 6:529–548, 1997.

Michael P. Wellman, Amy Greenwald, Peter Stone, and Peter R. Wurman. The 2001 trading agent competition. *Electronic Markets*, 13:4–12, 2003.

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.