Simultaneous Ascending Auctions (SAA): Strategies

December 19, 2007

In this appendix we describe bidding strategies we have designed to study simultaneous ascending auction games [Reeves et al., 2003, MacKie-Mason et al., 2004, Reeves et al., 2005, Osepayshvili et al., 2005, Wellman et al., 2007]. Throughout this appendix we use the concept of an *SAA environment*, which comprises an SAA mechanism over M goods, a set of N agents, and a probability distribution D over M-good value functions for each agent. We refer to this distribution as the distribution of the agents' preferences, and label the environment $D(M, N)$.

In the first five sections we describe the strategy families and the details of their implementation. Refer to Table 1 for the content of these sections. In sections 6 through 8 we report the strategy pools in the restricted games we studied in [Osepayshvili et al., 2005] and [Wellman et al., 2007].

1 Straightforward Bidding Strategy

A straightforward bidder (SB) takes a vector of *perceived prices* for the goods as given, and bids those prices for the bundle of goods that would maximize the agent's surplus if it were to win all of its bids at those prices. The perceived price of a good for an SB agent is equal to the current price of that good if the agent is winning the good, and to the current price plus the minimum price increment (set by the auction rules) if the agent is not winning the good. We describe details of SB in prior work [Reeves et al., 2005].

2 Sunk-Aware Strategies

A sunk-aware agent is a modification of the SB agent that bids as if the incremental cost for goods the agent is currently winning is somewhere on the interval of zero and the current bid price. To implement this modification, we introduced the sunk awareness parameter, $k \in [0, 1]$. If $k = 1$ the strategy is identical to straightforward bidding. At $k = 0$ the agent is fully sunk aware, bidding as if it would retain the goods it is currently winning with certainty.

Strategy Family and Notation	Family Parameter	Examples	Appendix Section
Straightforward bidder, SB	N/A	SB	1
Sunk-Aware Agent, SA(k)	Sunk-awareness parameter, k	SA(k) $k = \{0, 0.05, 0.1, \ldots 0.95\}$	2
Point Price Predictors, PP($\boldsymbol{\pi}^x$)	Beliefs about average final prices of the goods, x	PP($\boldsymbol{\pi}^{Zero}$) PP($\boldsymbol{\pi}^{\infty}$) PP($\boldsymbol{\pi}^{ECE}$) PP($\boldsymbol{\pi}^{EDCE}$) PP($\boldsymbol{\pi}^{SB}$) PP($\boldsymbol{\pi}^{SC}$)	3
Point Price Predictors with participation only, PP($\boldsymbol{\pi}^x$) w/ P.O.	Beliefs about average final prices of the goods, x	PP($\boldsymbol{\pi}^{\infty}$) w/ P.O. PP($\boldsymbol{\pi}^{Zero}$) w/ P.O. PP($\boldsymbol{\pi}^{ECE}$) w/ P.O. PP($\boldsymbol{\pi}^{EDCE}$) w/ P.O. PP($\boldsymbol{\pi}^{SB}$) w/ P.O. PP($\boldsymbol{\pi}^{SC}$) w/ P.O.	3
Price Distribution Predictors, PP(F^x) PP($G(\mu(x), \sigma(y))$) PP($F(\boldsymbol{\pi}^x)$)	Beliefs about final price distributions. Such beliefs are labeled by x or a pair (x, y)	PP(F^{Zero}) PP(F^{U}) PP(F^{SB}) PP(F^{CE}) PP($G(\mu(CE), \sigma(CE))$) PP($G(\mu(SB), \sigma(SB))$) PP($F(\boldsymbol{\pi}^{ECE})$) PP($F(\boldsymbol{\pi}^{EDCE})$) PP($F(\boldsymbol{\pi}^{SB})$) PP($F(\boldsymbol{\pi}^{SC})$)	3
Demand Reduction Agent, DR(κ)	Demand reduction parameter, κ	DR(κ) $\kappa = \{1, 2, \ldots 30, 32, 34, 36, 38, 40, 44, 48, 50, 52, 56, 60, 70, 80, 90, 100, 110, 120\}$	4
Own-Effect Price Predictor, OEPP($\boldsymbol{\pi}^x$)	Beliefs about own effect on final prices	OEPP($\boldsymbol{\pi}^{SB}$)	5

Table 1: Reference table.

Intermediate values are consistent with bidding as if the agent puts an intermediate probability on the likelihood of retaining the goods it is currently winning. We label sunk aware agents as SA(k). We describe the details of this strategy family in prior work [Reeves et al., 2005].

3 Price Predicting Strategies

Agents using price prediction strategies generate an initial, pre-auction belief about the final prices of the goods. For example, the agent can believe that the final prices will all equal zero or they are distributed normally or uniformly on some interval. The derivation of most beliefs we employ involves Monte Carlo sampling. To obtain a belief for a particular distribution of agents' preferences, we simulate a large number of game instances with agents drawn from the preference distribution.

In most of our reported results we used the uniform and exponential distributions of number of goods demanded by the agent. In our earlier studies we also considered agents who demanded a fixed number of goods. We refer to such a preference distribution as a fixed distribution.

We consider two families of price predictors (see Osepayshvili et al. [2005]). The *point prediction* strategy family has point beliefs about the final prices that will be realized for each good. The point price predicting strategy is thus parameterized by its vector of point beliefs. We label such beliefs $\boldsymbol{\pi}(\boldsymbol{\Omega_0}, \phi)$.[1] Any price vector with the number of elements equal to the number of goods can represent initial beliefs. Point beliefs based on sampling are obtained by averaging across final prices or demands in the simulated games.

The point prediction family includes a sub-class of strategies with participation-only-prediction. The idea behind this strategy is that it ignores its predictions at some stages of decision making (unlike the full point predictor, which always relies on the predicted vector). We describe this strategy variant in detail elsewhere [MacKie-Mason et al., 2004].

The *distribution prediction* strategy family has beliefs about the final price distributions. Let $F \equiv F(\boldsymbol{\Omega_0}; \phi)$ denote a joint cumulative distribution function over final prices, representing the agent's pre-auction belief. We assume that prices are bounded above by a known constant, V. Thus, F associates cumulative probabilities with price vectors in $\{0, \ldots, V\}^M$. For simplicity we use only the information contained in the vector of marginal distributions, $(F_1, \ldots, F_M)$, as if the final prices are independent across goods.

For technical reasons we require that initial beliefs are such that for each good all prices in $\{0, \ldots, V\}$ have a positive probability of occurring. To ensure that this requirement is satisfied for initial beliefs obtained from empirical samples, we modify the latter before constructing beliefs. In particular, we add to the empirically obtained sample of prices a sample of $(V + 1)$ prices from 0 to V. Suppose the empirical sample consists of n games. The probability that the

[1] Here and later in the appendix ϕ refers to the empty set of bid information available pre-auction.

good will have final price $p \in \{0, \ldots, V\}$ is then

$$\Pr(p) = \frac{n_p + 1}{n + (V + 1)}, \tag{1}$$

where n_p is the number of times the final price equaled p in the original sample.

If n is high relative to V, the effect of adding an artificial sample to the empirical one is negligible. In the restricted games we studied $V = 50$ and $n \geq 340000$.

In the following sections we describe some examples of beliefs. We denote a specific point price prediction strategy by PP($\boldsymbol{\pi}^x$), where x labels particular initial point beliefs, $\boldsymbol{\pi}(\boldsymbol{\Omega_0}, \phi)$. We denote the strategy of bidding based on a particular distribution predictor by PP(F^x), where x labels various initial beliefs about final price distributions, $F(\boldsymbol{\Omega_0}; \phi)$. If the initial beliefs x are based on Monte Carlo sampling, we write x_u and x_e to distinguish between beliefs obtained using draws from uniform and exponential preference distributions.[2]

3.1 Zero Beliefs

To construct zero beliefs for the distribution predictor, we assume that we have a sample of $n = 1000000$ zero final prices. After we add another artificial sample of $(V + 1) = 51$ prices (see equation (1)), the probability that the final price of a good will be zero becomes $1000000/1000051 = 0.99995$. The probability that the final price of this good will equal $p \in [1, V]$ becomes $1/1000051 < 10^{-6}$. Note, however, that as soon as the ask price exceeds zero, the distribution predictor reconditions its beliefs based on that information. So once all the ask prices are above zero it bids identically to distribution predictor having a uniform distribution for final prices.

For the point price predictor, zero beliefs are simply a vector of zeros. The point price predictor with zero beliefs is equivalent to SB (see Section 1).

We denote the point and distribution predictors using zero beliefs by PP($\boldsymbol{\pi}^{Zero}$) and PP(F^{Zero}) respectively.

3.2 Infinite Point Beliefs

Infinite point beliefs are defined only for the point price predictor. We label such a predictor PP($\boldsymbol{\pi}^\infty$). We define this strategy so that it reverts to SB if the agent has single-unit preference (see Osepayshvili et al. [2005]). Since there is no exposure problem, the agent will bid if and only if it has a positive value for exactly one of the goods (despite the expectation that the price will be infinite).

[2] In our results [MacKie-Mason et al., 2004, Osepayshvili et al., 2005] all agents have beliefs derived for the preference distribution of their environment. We suppress the subscripts of beliefs to simplify the notation. Thus, if we consider a uniform environment, x refers to initial beliefs based on samples from uniformly distributed types. If we consider an exponential environment, x refers to initial beliefs based on samples from exponentially distributed types. One exception is the 53-strategy game that we constructed for the 5×5 uniform environment [Osepayshvili et al., 2005]. In this game some predictors had beliefs based on the exponential distribution. The 53 strategies are described in Section 6.

The average performance of this strategy across a Monte Carlo sample provides a useful performance benchmark for price predicting strategies.

3.3 Uniform Distribution Beliefs

The uniform distribution beliefs are defined only for the distribution predictor, which we denote PP(F^U). According to the uniform beliefs, probability that the final price of a good will equal $p \in [0, V]$ is $1/51 = 0.0196$.

3.4 SB Beliefs

SB beliefs are based on final prices in a sample of games in which players used the SB strategy (see Section 1). The point SB beliefs are a vector of average final prices of the goods available. The distribution SB beliefs are a vector of marginal price distributions computed according to equation (1). Our sample size is $n = 1000000$ games, and the upper bound on prices is $V = 50$ for all goods.

We denote the point predictors with SB beliefs derived using samples from the uniform and exponential distributions by PP($\boldsymbol{\pi}^{SB_u}$) and PP($\boldsymbol{\pi}^{SB_e}$) respectively. Similarly, the distribution predictors with SB beliefs are labeled PP(F^{SB_u}) and PP(F^{SB_e}).

3.5 Competitive Equilibrium Beliefs

Suppose that the final prices form a competitive (or Walrasian) equilibrium (CE) in the SAA game. This is guaranteed, for example, when SB (see section 1) all demand only single goods, but is not true in general. However, in our experience, the final prices are generally not too far from CE prices. Therefore, we calculate the Walrasian equilibrium for an SAA environment and use the resulting prices to create initial beliefs.

Let SE be an SAA environment. To find the vector of *competitive equilibrium price distributions* for SE, F^{CE}, we randomly generate many ($n = 25000$) game instances with agents drawn from the preference distribution of SE, and use tatonnement to solve for the equilibrium prices in each. The CE distribution beliefs are computed based on the sample of the equilibrium prices according to the equation (1). We label the CE distribution predictor as PP(F^{CE_u}) if the preference distribution of the sample is uniform, and as PP(F^{CE_e}) if it is exponential.

We calculate the *expected competitive equilibrium* beliefs, $\boldsymbol{\pi}^{ECE}$, by averaging across the prices in the sample. We calculate the *expected demand competitive equilibrium* beliefs, $\boldsymbol{\pi}^{EDCE}$, by calculating the expected demand function for each of the game instances, and then solving for the competitive equilibrium based on the average demands. We label the corresponding bid strategies as PP($\boldsymbol{\pi}^{ECE_u}$) and PP($\boldsymbol{\pi}^{EDCE_u}$) if the preference distribution of the sample is uniform, and as PP($\boldsymbol{\pi}^{ECE_e}$) and PP($\boldsymbol{\pi}^{EDCE_e}$) if it is exponential.[3]

[3] The prices to which tatonnement converges are sensitive to the choice of initial prices and

3.6 Self-Confirming Prices

The self-confirming (SC) distribution beliefs are what we call self-confirming marginal price distributions. 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.

The SC point beliefs are what we call *self-confirming point predictions*, which are defined as a vector of point predictions that on average are correct, if all agents use point price prediction ($\boldsymbol{\pi}$). Note that the mean of a SC distribution may be different from SC point predictions for the same environment.

The general idea behind the derivation algorithms if both types of self-confirming predictions is as follows. Given an SAA environment, we derive self-confirming predictions through an iterative simulation process. Starting from an arbitrary prediction, we run many instances (n) of an SAA environment (sampling from the given preference distributions) with all agents playing the same predicting strategy (either point prediction or distribution prediction). We record the resulting prices from each instance, and create new beliefs for the predictors. The new point beliefs are a vector of average final prices from the first iteration. The new distribution beliefs are the sample distribution. We repeat the process using the new beliefs for each new iteration. If it ever reaches an approximate fixed point, then we have statistically identified approximate self-confirming predictions for this environment. For more details on SC point and distribution predictions, see MacKie-Mason et al. [2004] and Osepayshvili et al. [2005] respectively.

For the environment we investigate, we could find both SC point predictions and distributions. In our experiments $n = 500000$ for SC point predictions and $n = 1000000$ for SC distributions; the initial beliefs used in the first iteration are zero beliefs (see 3.1), although our results do not appear sensitive to this.

3.7 Gaussian Distribution Beliefs

The Gaussian distribution beliefs are defined only for the distribution predictor. Suppose we know the expected prices $\boldsymbol{\mu}$ and the standard deviations $\boldsymbol{\sigma}$ for all the goods. Then we can approximate the final price distribution of good m with a Gaussian centered on μ_m with the restriction that prices $p \in [0, V]$. To implement Gaussian beliefs we draw a random sample of size $n = 1000000$ from $N(\mu_m, \sigma_m)$ for each good m and discard prices outside $[0, V]$.[4] Then we compute the final price probabilities according to equation (1). We denote distribution predictors with such beliefs by $G(x, y)$, where x and y label the expected prices and standard deviations respectively. For example, agent $\mathrm{PP}(G(\mu(CE_u), \sigma(CE_u)))$ has Gaussian beliefs created using

other parameters of the algorithm. In Section 6 we list all the price vectors we obtained for the 5×5 uniform environment. At the time of writing this appendix, we have not derived CE beliefs for alternative environments.

[4] One drawback of this approach is that truncating the tails shifts the means of the distribution.

the means and standard deviations of the distribution vector F^{CE_u}; agent $\text{PP}(G(\boldsymbol{\pi}^{EDCE_u}, \sigma(CE_u)))$ has Gaussian beliefs created using $\boldsymbol{\pi}^{EDCE_u}$ as the means vector and the standard deviations of the distribution vector F^{CE_u}.

3.8 Degenerate Distribution Beliefs

The degenerate distribution beliefs are defined only for the distribution predictor. They are degenerate in the sense that according to such beliefs each good has a deterministic final price. Let $\boldsymbol{\pi}$ be a vector of prices. For every good m, assign probability 1 to m'th element of $\boldsymbol{\pi}$ and probability 0 to all other prices in $[0, V]$. We denote the distribution predictor with such beliefs by $\text{PP}(F(\boldsymbol{\pi}))$. If we round the vector $\boldsymbol{\pi}$ before generating the distributions, we mark the vector by a prime in the strategy name: $\text{PP}(F(\boldsymbol{\pi}'))$. For example, $\text{PP}(F(\boldsymbol{\pi}^{SB_u}))$ has beliefs that the final prices will equal $\boldsymbol{\pi}^{SB_u}$ with probability 1; $\text{PP}(F(\boldsymbol{\pi}^{ECE'_u}))$ has beliefs that the final prices will equal the rounded elements of $\boldsymbol{\pi}^{ECE_u}$ with probability 1. The former beliefs match better the information that the corresponding point predictor has. The latter beliefs can be justified by the fact that prices are integers in our model.

4 Demand Reduction Strategy

The demand reduction strategy family, $\text{DR}(\kappa)$, has been designed for homogeneous goods environments (see Wellman et al. [2007]). In such an environment, each auction sells one unit of a homogeneous indivisible good, and the bidders' marginal value for units of this good is weakly decreasing. We implemented such preferences by randomly drawing marginal values v_k for the kth good from $U[0, v_{k-1}]$, with v_0 a uniform upper bound on the marginal value of one unit equal to 127. The parameter $\kappa \in 0, \ldots, 127$ defines the degree of the agent's demand reduction. Larger κ implies that the agent bids on fewer units. When $\kappa = 0$, the agent's bidding behavior is equivalent to SB (see Section 1). In the other extreme case when $\kappa = 127$, the agent never bids at all. For more details on the strategy, see Wellman et al. [2007].

5 Own-Effect Price Predictors

Own-effect price predictor, $\text{OEPP}(\boldsymbol{\pi}^x)$, is an extension of the price predictor(see Section 3 in the Appendix), designed for homogeneous goods environments (see Wellman et al. [2007]). Its belief is an $m \times m$ matrix $(\pi_{ik}(\boldsymbol{B}))$, where $\pi_{ik}(\boldsymbol{B})$ is a predicted final price of good i given that the agent tries to win k goods and its information state at the current round is $\boldsymbol{B}$.

As we argue in Wellman et al. [2007], there is no obvious reason why an agent should believe that the final price of a homogeneous good on one auction will be higher than the price on another auction. Due to this consideration, we construct the initial price prediction to be equal across auctions: $\pi_{ik}(\varnothing) = \pi_{jk}(\varnothing)$ for all i and j for all purchase sizes k. In other words, the elements in a column are identical in the agent's initial prediction matrix. We label the initial prediction matrix of *predicted own-effect prices* by π^x, in which the subscript x labels particular initial predictions.

Any $m \times m$ matrix can potentially be the agent's beliefs. We have studied *self-confirming own price effects*, i.e. beliefs that on average are correct.

5.1 Self-Confirming Own-Effect Prices

Let Γ be an instance of an SAA game with homogeneous goods. A vector $\boldsymbol{\pi}$ is *self-confirming own-effect price vector for* Γ, if for all $k \in \{1, \dots, m\}$, the final price is equal to the expectation (over the type distribution) of the final price when one agent tries to win k goods and all the other agents follow OEPP($\boldsymbol{\pi}$).

We denote the self-confirming own-effect price vector by $\boldsymbol{\pi}^{SC}$ and the self-confirming own-effect price prediction strategy by OEPP($\boldsymbol{\pi}^{SC}$).

To find approximate self-confirming own-effect prices trough an iterative procedure. First, we initialize the price effect predictors' beliefs with an arbitrary matrix and simulate many games with a profile in which one agent (the *explorer*) ignores its preferences and tries to win a single unit of the good, while the others follow OEPP. When average prices obtained by these agents are determined, we replace the first element in the price effects vector with the average price, re-program the explorer to win two goods and repeat. After the second batch of simulations, we replace the second element of the price effects vector with the average price and increase the explorer's bundle size s by one or, if $s + 1 > m$, return to $s = 1$. When this process reaches a fixed point, we have the self-confirming vector of own-effect prices, $\boldsymbol{\pi}^{SC}$.

We implemented OEPP($\boldsymbol{\pi}^{SC}$) for a homogeneous environment with 5 goods and 5 agents (see Wellman et al. [2007]). We set the initial beliefs to a matrix of zeros, and simulated 10 thousand games for every bundle size of the explorer. The explorer's bundle size was changed 500 times, i.e., each bundle size was updated 100 times. We created 10 approximate self-confirming vectors through this procedure and used their average as $\boldsymbol{\pi}^{SC}$. The approximate self-confirming vectors and their average are presented in Table 2.

6 53-Strategy Game For 5×5 Uniform Complementary Environment

In this section we describe 53 strategies that we constructed for the 5×5 uniform environment (see Osepayshvili et al. [2005]). In Table 3 we list all the strategies, relevant sections of this appendix and all our publications in which the strategies were mentioned. Table 4 presents all initial point beliefs for the price predicting strategies we derived for 5×5 uniform environment, as well as relevant sections of this appendix.

1	2	3	4	5
19.0538	39.994	60.934	80.4613	101.15
19.7819	39.6585	60.9936	80.3963	101.755
19.3541	40.2281	60.9409	80.4295	102.034
19.2679	39.8432	60.8178	80.2026	101.359
18.5337	39.5263	60.5431	80.4075	101.501
19.5411	40.2917	61.0213	80.5834	101.373
18.9921	40.0143	61.2818	80.377	101.792
18.2261	39.8009	61.2354	80.4773	101.684
18.73	39.5478	61.2418	80.572	101.787
18.4487	39.7332	61.4828	80.5034	101.204
18.99294	39.8638	61.04925	80.44103	101.5639

Table 2: Approximate self-confirming own-effect price vectors and their average. The columns are the target purchase sizes.

7 Strategies For Alternative Complementary Environments

Table 5 describes a pool of 27 strategies, which we searched for for a profitable deviation from playing PP(F^{SC}) when all the other agents play PP(F^{SC}). Table 6 describes 7-strategy games (7-cliques) for each of the alternative environments we considered in Osepayshvili et al. [2005]. Some of the strategies are price predictors whose initial beliefs are derived using Monte Carlo sampling. Such beliefs are therefore parameterized by the underlying preference distribution. However, we suppress the preference distribution labels in the tables, because we derived these beliefs only for the environment in which the corresponding predicting strategy was used, and therefore there is no ambiguity about how the beliefs were derived.

St. Family	#	Strategy Notation	Sec.	Publication
N/A	1	SB	1	Reeves et al. [2003] MacKie-Mason et al. [2004] Reeves et al. [2005] Osepayshvili et al. [2005]
Sunk- Aware	20	SA(k) $k = 0, 0.05, 0.1, \ldots 0.95$	2	Reeves et al. [2003] MacKie-Mason et al. [2004] Reeves et al. [2005] Osepayshvili et al. [2005]
Point Price Predictors	13	PP($\boldsymbol{\pi}^{\infty}$), PP($\boldsymbol{\pi}^{SB_u}$), PP($\boldsymbol{\pi}^{SB_u}$) w/ P.O., PP($\boldsymbol{\pi}^{ECE_u}$), PP($\boldsymbol{\pi}^{ECE_u^*}$), PP($\boldsymbol{\pi}^{ECE_u^{**}}$), PP($\boldsymbol{\pi}^{ECE_e}$), PP($\boldsymbol{\pi}^{ECE_e^*}$), PP($\boldsymbol{\pi}^{EDCE_u}$), PP($\boldsymbol{\pi}^{EDCE_u^*}$), PP($\boldsymbol{\pi}^{EDCE_e}$), PP($\boldsymbol{\pi}^{EDCE_e^*}$), PP($\boldsymbol{\pi}^{SC_u}$)	3	Reeves et al. [2005] Osepayshvili et al. [2005]
Price Distribution Predictors	19	PP(F^{Zero}),PP(F^U) PP(F^{SB_u}),PP(F^{SC_u}), PP(F^{CE_u}), PP($G(\mu(CE_u), \sigma(CE_u))$), PP($G(\mu(SB_u), \sigma(SB_u))$), PP($G(\boldsymbol{\pi}^{EDCE_u}, \sigma(CE_u))$), PP($G(\boldsymbol{\pi}^{EDCE_u}, \sigma(SB_u))$), PP($G(\boldsymbol{\pi}^{SC_u}, \sigma(CE_u))$), PP($G(\boldsymbol{\pi}^{SC_u}, \sigma(SB_u))$), PP($F(\boldsymbol{\pi}^{ECE_u})$), PP($F(\boldsymbol{\pi}^{ECE'_u})$), PP($F(\boldsymbol{\pi}^{EDCE_u})$), PP($F(\boldsymbol{\pi}^{EDCE'_u})$), PP($F(\boldsymbol{\pi}^{SB_u})$), PP($F(\boldsymbol{\pi}^{SB'_u})$) PP($F(\boldsymbol{\pi}^{SC_u})$), PP($F(\boldsymbol{\pi}^{SC'_u})$)	3	Osepayshvili et al. [2005]

Table 3: 53 strategies for the largest 5×5 uniform environment game. The strategies are listed in column 3, and their strategy families are given in column 1. P.O. refers to participation only prediction. Different point beliefs created by the same (non-deterministic) algorithm are marked by asterisks. In column 2 we report the total number of strategies from a particular family represented in the 53-strategy game. The relevant Appendix sections are reported in column 4. Our publications, in which the strategies were mentioned, are listed in column 5.

Beliefs/Good	1	2	3	4	5	Appendix Section
	Initial Point Beliefs					
$\boldsymbol{\pi}^{\infty}$	1000	1000	1000	1000	1000	3.2
$\boldsymbol{\pi}^{SB_u}$	14.8	10.7	7.6	4.6	1.9	3.4
$\boldsymbol{\pi}^{SC_u}$	13.0	8.7	5.4	3.0	1.2	3.6
$\boldsymbol{\pi}^{ECE_u}$	16.6	10.8	6.5	3.1	0.7	3.5
$\boldsymbol{\pi}^{ECE_u^*}$	16.5	10.7	6.4	3.1	0.8	3.5
$\boldsymbol{\pi}^{ECE_u^{**}}$	26.0	14.2	6.9	2.5	0.3	3.5
$\boldsymbol{\pi}^{ECE_e}$	6.0	4.1	1.8	0.6	0.1	3.5
$\boldsymbol{\pi}^{ECE_e^*}$	30.5	11.9	6.0	2.7	0.4	3.5
$\boldsymbol{\pi}^{EDCE_u}$	20.0	12.0	8.0	2.0	0.0	3.5
$\boldsymbol{\pi}^{EDCE_u^*}$	20.8	11.4	8.2	1.8	0.0	3.5
$\boldsymbol{\pi}^{EDCE_e}$	25.0	10.0	5.1	0.9	0.0	3.5
$\boldsymbol{\pi}^{EDCE_e^*}$	24.5	10.5	5.5	1.5	0.0	3.5

Table 4: Initial beliefs (rounded to one decimal point) for 5×5 uniform environment. The notation for the beliefs is presented in column 1. The vectors of point beliefs are presented in column 2. The goods are numbered from 1 through 5. The monotonicity of the prices is due to the specifics of the scheduling games (see MacKie-Mason et al. [2004] and Reeves et al. [2005]) we considered in our experiments. For information about price predicting strategies, see Section 3. The subsections of Section 3 that are most relevant to a particular belief are presented in column 3.

Strategy Family	#	Strategy Notation	Section
N/A	1	SB	1
Sunk-Aware	20	SA(k) $k = \{0, 0.05, 0.1, \ldots 0.95\}$	2
Point Price Predictors	3	PP($\boldsymbol{\pi}^{\infty}$), PP($\boldsymbol{\pi}^{SB}$), PP($\boldsymbol{\pi}^{SC}$)	3
Price Distribution Predictors	3	PP(F^U), PP(F^{SB}), PP(F^{SC})	3

Table 5: Pool of 27 deviators for alternative environments. The strategies are listed in column 3, and their strategy families are given in column 1. Column 2 gives then total number of strategies from a particular family represented in the 27-strategy pool. The relevant Appendix sections are listed in column 4.

Environment	Strategies
$E(3,3)$	PP(F^{SC}), PP(F^{SB}), SA(0.6), SA(0.7), SA(0.75), SA(0.8), SA(0.85)
$E(3,5)$	PP(F^{SC}), PP(F^{SB}), PP(F^{U}), PP($\boldsymbol{\pi}^{SC}$), PP($\boldsymbol{\pi}^{SB}$), PP($\boldsymbol{\pi}^{\infty}$), SA(0.85)
$E(3,8)$	PP(F^{SC}), PP(F^{SB}), PP(F^{U}), PP($\boldsymbol{\pi}^{SC}$), PP($\boldsymbol{\pi}^{SB}$), PP($\boldsymbol{\pi}^{\infty}$), SA(0.85)
$E(5,3)$	PP(F^{SC}), PP(F^{SB}), PP(F^{U}), PP($\boldsymbol{\pi}^{SC}$), PP($\boldsymbol{\pi}^{SB}$), PP($\boldsymbol{\pi}^{\infty}$), SB
$E(5,5)$	PP(F^{SC}), PP(F^{SB}), PP(F^{U}), PP($\boldsymbol{\pi}^{SC}$), PP($\boldsymbol{\pi}^{SB}$), PP($\boldsymbol{\pi}^{\infty}$), SA(0.35)
$E(5,8)$	PP(F^{SC}), PP(F^{SB}), PP(F^{U}), PP($\boldsymbol{\pi}^{SC}$), PP($\boldsymbol{\pi}^{SB}$), PP($\boldsymbol{\pi}^{\infty}$), SA(0.35)
$E(7,3)$	PP(F^{SC}), PP(F^{SB}), SA(0.55), SA(0.65), SA(0.7), SA(0.75), SA(0.8)
$E(7,6)$	PP(F^{SC}), PP(F^{SB}), PP(F^{U}), PP($\boldsymbol{\pi}^{SC}$), PP($\boldsymbol{\pi}^{SB}$), PP($\boldsymbol{\pi}^{\infty}$), SB
$U(3,3)$	PP(F^{SC}), PP(F^{SB}), SA(0.65), SA(0.7), SA(0.75), SA(0.8), SA(0.85)
$U(3,5)$	PP(F^{SC}), PP(F^{SB}), PP($\boldsymbol{\pi}^{SC}$), PP($\boldsymbol{\pi}^{SB}$), SA(0.7), SA(0.75), SA(0.85)
$U(3,8)$	PP(F^{SC}), PP(F^{SB}), PP(F^{U}), PP($\boldsymbol{\pi}^{SC}$), PP($\boldsymbol{\pi}^{SB}$), PP($\boldsymbol{\pi}^{\infty}$), SA(0.9)
$U(5,3)$	PP(F^{SC}), PP(F^{SB}), SA(0.6), SA(0.65), SA(0.7), SA(0.75), SA(0.8)
$U(5,8)$	PP(F^{SC}), PP(F^{SB}), PP(F^{U}), PP($\boldsymbol{\pi}^{SC}$), PP($\boldsymbol{\pi}^{SB}$), PP($\boldsymbol{\pi}^{\infty}$), SA(0.9)
$U(7,3)$	PP(F^{SC}), PP(F^{SB}), SA(0.55), SA(0.65), SA(0.7), SA(0.75), SA(0.8)
$U(7,6)$	PP(F^{SC}), PP(F^{SB}), PP($\boldsymbol{\pi}^{SC}$), PP($\boldsymbol{\pi}^{SB}$), SA(0.75), SA(0.8), SA(0.9)
$U(7,8)$	PP(F^{SC}), PP(F^{SB}), PP(F^{U}), PP($\boldsymbol{\pi}^{SC}$), PP($\boldsymbol{\pi}^{SB}$), PP($\boldsymbol{\pi}^{\infty}$), SB

Table 6: 7-strategy games for alternative environments. $E(M,N)$ and $U(M,N)$ refer to the environments with exponential and uniform preference distribution respectively, M goods and N agents.

8 51-Strategy Game for 5×5 Uniform Homogeneous Environment

In Table 7 we report 51 strategies that we constucted for the 5×5 homogeneous goods environment (Wellman et al. [2007]), relevant sections of this appendix and all our publications in which the strategies have been mentioned.

St. Family	#	Strategy Notation	Sec.	Publication
N/A	1	SB	1	Reeves et al. [2003] MacKie-Mason et al. [2004] Reeves et al. [2005] Osepayshvili et al. [2005] Wellman et al. [2007]
Sunk-Aware	1	SA(k) $k = 0.5$	2	Reeves et al. [2003] MacKie-Mason et al. [2004] Reeves et al. [2005] Osepayshvili et al. [2005] Wellman et al. [2007]
Price Distribution Predictor	1	PP(F^{SB_u})	3	Osepayshvili et al. [2005] Wellman et al. [2007]
Demand Reduction Agent	47	DR(κ) $\kappa = \{1, 2, \ldots 30, 32, 34, 36, 38, 40, 44, 48, 50, 52, 56, 60, 70, 80, 90, 100, 110, 120\}$	4	Wellman et al. [2007]
Own-Effect Price Predictor	1	OEPP($\boldsymbol{\pi}^{SC}$)	5	Wellman et al. [2007]

Table 7: 51 strategies for the 5×5 homogeneous goods environment game. The strategies are listed in column 3, and their strategy families are given in column 1. In column 2 we report the total number of strategies from a particular family represented in the 51-strategy game. The relevant Appendix sections are reported in column 4. Our publications, in which the strategies were mentioned, are listed in column 5.

References

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*, Whistler, BC, 2004.

Anna Osepayshvili, Michael P. Wellman, Daniel M. Reeves, and Jeffrey K. MacKie-Mason. Self-confirming price prediction for bidding in simultaneous ascending auctions. 21st Conference on Uncertainty in Artificial Intelligence (UAI 2005), July 2005.

Daniel M. Reeves, Michael P. Wellman, Jeffrey K. MacKie-Mason, and Sowmya Swaminathan. Exploring bidding strategies for market-based scheduling. San Diego, June 2003. Fourth ACM Conference on Electronic Commerce.

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

Osepayshvili. Exploring bidding strategies for market-based scheduling. *Decision Support Systems*, 39(1), March 2005.

Michael P. Wellman, Anna Osepayshvili, Jeffrey K. MacKie-Mason, and Daniel M. Reeves. Bidding strategies for simultaneous ascending auctions, 2007. Working paper.