Bidding Strategies for Simultaneous Ascending Auctions: Appendix

December 11, 2008

In Section 1 of this appendix, we provide a complete parameter specification for all bidding strategies we have designed for simultaneous ascending auctions (SAAs) (Reeves et al. [2003], MacKie-Mason et al. [2004], Reeves et al. [2005][1], Osepayshvili et al. [2005], Wellman et al. [2008]). In Section 2, we describe some restricted SAA games we have analyzed.

The formal specification of the SAA game includes a number of agents, n, a number of goods, m, a type distribution that yields valuation functions, v_j for the agents $j \in \{1, \dots, n\}$, and a specification of the SAA mechanism rules. In general, the SAA mechanism comprises m separate auctions, one for each good $i \in \{1, \dots, m\}$, that operate over multiple rounds of bidding. In the generic SAA version we study, bidding is synchronized so that in each round each agent submits a bid in every auction in which it chooses to bid. At any given time, the *bid price* on good i is β_i, defined to be the highest bid b_i received thus far, or zero if there have been no bids. The bid price along with the current winner in every auction is announced at the beginning of each new round. To be admissible, a new bid must meet the *ask price*, i.e., the bid price plus a bid increment (which we take to be one w.l.o.g., allowing for scaling of the agent values): $b_i^{new} \geq \beta_i + 1$. If an auction receives multiple admissible bids in a given round, it admits the highest, breaking ties randomly. An auction is *quiescent* when a round passes with no new admissible bids, i.e., the new bid prices $\boldsymbol{\beta}^{new} = \boldsymbol{\beta}$ which become the final prices $\boldsymbol{p}$. When every auction is simultaneously quiescent they all close, allocating their respective goods per the last admitted bids.

To analyze strategies for SAAs, we employ *empirical game-theoretic methodology* ([MacKie-Mason and Wellman, 2005] and [Wellman, 2006]), which involves simulation of many SAA game instances in a variety of *environments*. Particular environments are defined by specifying the number m of goods, the number n of bidding agents, and a preference model of agents.

Some strategies described in this document are based on environment-specific price predictions. In Section 1 we provide a description of the prediction methods for these strategies. An index of strategies and corresponding prediction methods is given in Table . For each strategy, we provide the notation, the strategy parameter (if any), the section of the appendix in which the strategy is described, as well as references to our publications in which the strategy appears in an SAA game.[2]

1 Bidding Strategies

All strategies described in this document belong to the strategy class that we call *perceived-price bidders*, which is parametrized by a *perceived-price function* $\rho \colon \boldsymbol{\mathcal{B}} \to \mathbb{Z}_*^m$ which maps the agent's information state, $\boldsymbol{B}$, to a (nonnegative, integer) perceived-price vector, $\rho(\boldsymbol{B})$.[3]

A perceived-price bidder computes the subset of goods

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

breaking ties in favor of smaller subsets and lower-numbered goods.Then, given X^*, the agent bids $b_i = \beta_i + 1$ (the ask price) for the $i \in X^*$ that it is not already winning.

In the following sections we provide details on how we constructed the perceived-price function for each bidding strategy. Most of our bidding strategies are well defined for any type distribution. Two exceptions are the demand-reduction strategy (Section 1.4) and the own-effect price predictor (Section 1.5) defined only for homogeneous-good environments. See the corresponding publications for motivation, analysis and discussion of the strategies (see the strategy index in Table).

[1] Reeves et al. [2005] is a more recent and extended version of Reeves et al. [2003].

[2] There is a typo in the description of the strategy space in Osepayshvili et al. [2005]. In Section 5.1 "Environments and strategy space", we say that our strategies were drawn from three strategy families: SB, point predictor, and distribution predictor. In fact, the list should also include a fourth strategy family, the sunk-aware bidder. The strategy list is fully specified in Table .

[3] See, for example, Wellman et al. [2008] for a formal definition of information state.

Strategy & notation	Strategy parameter	Example	Appendix Section	Publications
Straightforward bidder, SB	N/A	SB	1.1	*Reeves et al. [2003]* MacKie-Mason et al. [2004] *Reeves et al. [2005]* Osepayshvili et al. [2005] Wellman et al. [2008]
Sunk-aware agent, SA(k)	Sunk-awareness parameter, k	SA(k) $k = \{0, 0.05, 0.1, \ldots 0.95\}$	1.2	*Reeves et al. [2003]* MacKie-Mason et al. [2004] *Reeves et al. [2005]* Osepayshvili et al. [2005] Wellman et al. [2008]
Point price predictor, PP($\boldsymbol{\pi}^x$)	Predictions about average final prices of the goods, $\boldsymbol{\pi}^x$, where x labels method of generating prices	PP($\boldsymbol{\pi}^{Zero}$) PP($\boldsymbol{\pi}^{\infty}$) PP($\boldsymbol{\pi}^{EPE}$) PP($\boldsymbol{\pi}^{EDPE}$) PP($\boldsymbol{\pi}^{SB}$) PP($\boldsymbol{\pi}^{SC}$)	1.3	*MacKie-Mason et al. [2004]* Osepayshvili et al. [2005] Wellman et al. [2008]
Point price predictor with participation only, PP($\boldsymbol{\pi}^x$) w/ P.O.	Predictions about average final prices of the goods, $\boldsymbol{\pi}^x$, where x labels method of generating prices	PP($\boldsymbol{\pi}^{\infty}$) w/ P.O. PP($\boldsymbol{\pi}^{Zero}$) w/ P.O. PP($\boldsymbol{\pi}^{EPE}$) w/ P.O. PP($\boldsymbol{\pi}^{EDPE}$) w/ P.O. PP($\boldsymbol{\pi}^{SB}$) w/ P.O. PP($\boldsymbol{\pi}^{SC}$) w/ P.O.	1.3	*MacKie-Mason et al. [2004]* Osepayshvili et al. [2005] Wellman et al. [2008]
Price distribution predictor, PP(F^x) PP($G(\mu(x), \sigma(y))$) PP($F(\boldsymbol{\pi}^x)$)	Predictions about marginal final-price distributions, F, where F is labeled by method of generating prices	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($F(\boldsymbol{\pi}^{EPE})$) PP($F(\boldsymbol{\pi}^{EDPE})$) PP($F(\boldsymbol{\pi}^{SB})$) PP($F(\boldsymbol{\pi}^{SC})$)	1.3	*Osepayshvili et al. [2005]* *Wellman et al. [2008]*
Demand- reduction agent, DR(κ)	Demand-reduction parameter, κ	DR(κ) $\kappa = \{1, 2, \ldots 30, 32, 34,$ $36, 38, 40, 44, 48, 50, 52, 56,$ $60, 70, 80, 90, 100, 110, 120\}$	1.4	*Wellman et al. [2008]*
Own-effect price predictor, OEPP($\boldsymbol{\pi}^x$)	Predictions about own effect on final prices, $\boldsymbol{\pi}^x$, where x labels method of generating prices	OEPP($\boldsymbol{\pi}^{SB}$)	1.5	*Wellman et al. [2008]*

Table 1: Strategy Index. In the last column we report publications in which the strategy appears in an SAA game. The most relevant publications are in italics. P.O. refers to *participation-only* prediction (see MacKie-Mason et al. [2004]).

1.1 Straightforward Bidding Strategy

A straightforward bidder (SB) sets $\rho(\boldsymbol{B})$ to *myopically perceived prices*, which equal the bid price for goods it was winning in the previous round and the ask price for the other goods:

$$\rho_i(\boldsymbol{B}) = \begin{cases} \beta_i & \text{if winning good } i \\ \beta_i + 1 & \text{otherwise,} \end{cases} \tag{1}$$

where $\boldsymbol{\beta}$ is the current bid prices.

1.2 Sunk-Aware Strategies

Our sunk-aware strategies generalize SB's method for choosing the perceived-price vector through the parameter $k \in [0, 1]$:

$$\rho_i(\boldsymbol{B}) = \begin{cases} \beta_i & \text{if the agent has a single-unit demand and is winning good } i \\ k\beta_i & \text{if the demand function is other than single-unit demand and the agent is winning good } i \\ \beta_i + 1 & \text{otherwise,} \end{cases}$$

where $\boldsymbol{\beta}$ is the current bid prices.

1.3 Prediction-Based Perceived-Price Strategies

A prediction-based strategy takes an environment-specific price prediction as input and uses it to construct a perceived-price function. Our prediction methods vary in the amount of information that becomes input to the strategy. A vector of average final prices over a history of auctions is an example of a *point prediction*. A distribution of such prices is an example of a *distribution prediction*. In our analysis, we use only the information contained in the marginal distributions of final prices of each good.

In each round of an SAA game, a price-predicting agent uses a set of rules to update its prediction based on price-quote information. Then the agent uses the updated prediction to generate its perceived prices for the current round. We describe our update rule and the method of constructing prediction-based perceived prices in MacKie-Mason et al. [2004], Osepayshvili et al. [2005], and Wellman et al. [2008]. In this appendix, we report how we constructed *pre-auction* predictions for each prediction-based strategy.

We employ Monte Carlo sampling to construct some of our initial predictions. To obtain a prediction for a particular distribution of agents' preferences, we simulate a large number of game instances with agents drawn from that preference distribution. Point price prediction based on Monte Carlo sampling is always an m-element vector obtained by averaging across final prices or demand draws. Distribution prediction based on Monte Carlo sampling is an m-vector of marginal price distributions, each computed according to Equation (2) below.

We have analyzed multiple preference models, including what we call uniform and exponential models for complementary and substitutable goods. See Reeves et al. [2005] for the scheduling problem with complementary values and Wellman et al. [2008], Section 7 for environments with substitutes. We denote such models (complementary/substitutable) $U(m, n)$ and (complementary/substitutable) $E(m, n)$ respectively, where m is the number of goods and n is the number of agents.

We assume that prices are bounded by some constant V. For technical reasons, we require that initial distribution predictions are such that for each good, all integer prices in $\{0, \ldots, V\}$ have a positive probability of occurring. To ensure that this requirement is satisfied for initial distribution predictions obtained from empirical samples, we add a one to the counts of all prices from 0 to V. This increases the sample size by $V+1$. In our empirical analysis, V is at most 0.3% of the size of the empirical sample, and therefore the effect on the shape of the probability distribution is negligible. In Table 2 we summarize for each environment the number of profiles, the minimum, maximum, and average number of game instances generated per strategy profile.

Let an empirical sample consist of N game instances. If we use this sample to generate an initial distribution price prediction, our marginal-probability estimate that a good will have final price $p \in \{0, \ldots, V\}$ is given by

$$\Pr(p) = \frac{N_p + 1}{N + (V + 1)}, \tag{2}$$

where N_p is the number of times the final price equals p in the original sample.

Environment	V	Num. Profiles	Min	Max	Average
Complementary $E(3,3)$	50	104	0.4	10.4	2.919
Complementary $E(3,5)$	50	462	0.6	0.65	0.643
Complementary $E(3,8)$	50	3023	0.37	10.44	0.453
Complementary $E(5,3)$	50	84	0.6	0.6	0.6
Complementary $E(5,5)$	50	462	0.6	0.6	0.6
Complementary $E(5,8)$	50	3023	0.5	9.61	0.601
Complementary $E(7,3)$	50	84	0.6	0.6	0.6
Complementary $E(7,6)$	50	924	0.6	0.6	0.6
Complementary $U(3,3)$	50	104	0.6	7.6	2.071
Complementary $U(3,5)$	50	462	0.4	6.4	0.72
Complementary $U(3,8)$	50	3023	0.41	6.44	0.48
Complementary $U(5,3)$	50	104	0.6	5.6	1.783
Complementary $U(5,5)$	50	4457	0.2	200.46	7.01
Complementary $U(5,8)$	50	3023	0.6	6.6	0.65
Complementary $U(7,3)$	50	104	0.6	6.61	2.049
Complementary $U(7,6)$	50	944	0.6	6.6	0.759
Complementary $U(7,8)$	50	3023	0.6	4.6	0.629
Substitutable $U(5,5)$	127	16995	0.04	34.485	0.986

Table 2: Minimum, maximum, and average number of game instances (in millions of samples) generated per strategy profile in complementary and substitutable environments. The number of profiles is a total over all restricted games analysed for the environment. These statistics describe the data accumulated over years up to 2008, including the simulations we describe in Wellman et al. [2008].

In the following sections we describe a few initial price predictions. We denote a specific point price-prediction strategy by PP($\boldsymbol{\pi}^x$), where x labels a particular initial point prediction. We denote the strategy based on a particular distribution predictor by PP(F^x), where x labels various initial predictions about final price distributions. If the initial prediction is based on Monte Carlo sampling, we write x_u and x_e to distinguish between predictions obtained using draws from uniform and exponential preference distributions.[4]

1.3.1 Zero Prediction

Zero *point* prediction is simply an m-element vector of zeros. We denote the strategy by PP($\boldsymbol{\pi}^{Zero}$). The bidding behaviour of PP($\boldsymbol{\pi}^{Zero}$) is identical to that of SB (defined in Section 1.1).

To construct a zero *distribution* prediction, we create for each good an artificial empirical sample in which the zero price occurs $N = 1,000,000$ times and the (integer) prices in $\{1, \ldots, V\}$ never occur. We then compute the marginal PDFs according to Equation (2). We denote the strategy by PP(F^{Zero}). Note that as soon as the ask price of a good exceeds zero, the updated belief for the good becomes uniform (see, for example, Wellman et al. [2008] for the update rule).

1.3.2 Infinite Point Prediction

By infinite prediction we mean a price prediction that is higher than the maximum price any agent is ever willing to pay given the agents' preference distribution. We implement it as an m-element price vector with each price equal to 1,000. We denote the strategy by PP($\boldsymbol{\pi}^{\infty}$). PP($\boldsymbol{\pi}^{\infty}$) serves as a useful performance benchmark for point price predictors. The agent bids if and only if it has single-unit demand, in which case its bidding is identical to that of SB.[5]

[4] In our results [MacKie-Mason et al., 2004, Osepayshvili et al., 2005], all agents have predictions derived for the preference distribution of their environment. We therefore suppress the subscripts of predictions to simplify the notation. Thus, if we consider a uniform environment, x refers to initial predictions based on samples from uniformly distributed types. If we consider an exponential environment, x refers to initial predictions 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 restricted game, some agents have predictions based on the exponential distribution. The 53 strategies are described in Section 2.1.

[5] Any perceived-price predictor reverts to SB if the agent has single-unit demand (see Osepayshvili et al. [2005] or Wellman et al. [2008]).

(Complementary) Environment	SB-based distribution mean (=SB-based point prediction)
$E(3,3)$	15.834215 6.742513 2.167808
$E(3,5)$	21.572768 11.496845 4.578351
$E(3,8)$	26.956882 15.574660 7.751081
$E(5,3)$	13.841055 5.874415 2.436862 1.033989 0.374938
$E(5,5)$	19.806997 10.267301 5.133310 2.357000 0.860561
$E(5,7)$	23.920096 13.139175 7.377028 3.701761 1.415634
$E(5,8)$	25.575856 14.218179 8.282323 4.331364 1.699206
$E(7,3)$	13.139589 5.471740 2.353370 1.141660 0.545570 0.250877 0.093792
$E(7,6)$	21.460556 11.246789 6.050605 3.339829 1.802822 0.897841 0.340184
$E(7,8)$	25.067160 13.645040 7.923657 4.673294 2.670973 1.400783 0.549261
$E(7,9)$	26.567146 14.595322 8.687265 5.263056 3.089004 1.654935 0.661331
$U(3,3)$	14.740242 8.353250 2.855183
$U(3,5)$	19.564241 12.732952 5.548269
$U(3,8)$	23.940851 16.026071 8.652534
$U(5,3)$	11.149274 7.59509 4.777189 2.613071 1.007106
$U(5,5)$	14.84235 10.70182143 7.55029571 4.63114857 1.89844518
$U(5,7)$	17.409806 12.449534 9.155785 6.120717 2.741584
$U(5,8)$	18.521817 13.104929 9.713525 6.696087 3.131634
$U(7,3)$	9.181489 6.707729 4.799898 3.369079 2.191272 1.276565 0.557992
$U(7,6)$	13.45445 10.019256 7.735138 5.977675 4.318716 2.69715 1.200078
$U(7,8)$	15.511484 11.287756 8.724902 6.904324 5.232313 3.458166 1.602136
$U(7,9)$	16.438418 11.824945 9.102629 7.244354 5.580744 3.780623 1.7924
$Ex1(2,2)$	15 14.50076

Table 3: SB-based price predictions in complementary environments. $Ex1(2,2)$ refers to the fixed preferences described in Wellman et al. [2008], Example 1.

1.3.3 Uniform Distribution Prediction

For each good, we create an artificial empirical sample in which all prices in $\{0,\ldots,V\}$ occur 20,000 times. We then compute the marginal PDFs according to Equation (2). We denote the uniform-distribution predictor by $\mathrm{PP}(F^U)$.

1.3.4 SB-Based (Baseline) Prediction

We simulate one million ($N = 1,000,000$) game instances in which all players follow the SB strategy (described in Section 1.1). The SB-based *point* prediction is an m-element vector of the average final prices. The SB-based *distribution* prediction an m-vector of marginal price distributions, each computed according to Equation (2). We denote the SB-based point prediction by $\mathrm{PP}(\boldsymbol{\pi}^{SB_x})$, where x labels the preference distribution of the SB-players: uniform (U) or exponential (E). Similarly, we denote the SB-based distribution predictor by $\mathrm{PP}(F^{SB_x})$. In our publications, we also refer to the SB-based prediction as baseline. In Table 3 we report the means of SB-based distribution predictions for a number of complementary environments. The means also represent SB-based point predictions (by definition).

1.3.5 Competitive Equilibrium Prediction

We repeatedly sample agents' preferences from the preference distribution (exponential or uniform) and apply Walrasian tatonnement to obtain a crude Monte Carlo estimate of the expected price equilibrium. The process is described in Wellman et al. [2008], Section 5.2. To construct a competitive-equilibrium distribution prediction, we apply Equation (2) to the sample of price-equilibrium estimates. We denote the strategy by $\mathrm{PP}(F^{EPE_x})$, where x labels the preference distribution. We also construct two types of point predictions, $\mathrm{PP}(\boldsymbol{\pi}^{EPE_x})$ and $\mathrm{PP}(\boldsymbol{\pi}^{EDPE_x})$. They differ in the order in which averaging and the tatonnement are applied. See Wellman et al. [2008] for details.

We have implemented this prediction method only for the $U(5,5)$ and $E(5,5)$ complementary environments. We found the prices to which tatonnement converges to be sensitive to the choice of initial prices and other parameters of the tatonnement algorithm. In Section 2.1 of this appendix, we provide all the

Environment	Prediction	Iterations	Games per iteration
$U(5,5)$a	Point	70	500,000
$U(5,5)$	Point	40	500,000
$U(5,5)$	Distribution	50	1 million
$E(7,9)$	Point	87	1 million
$Ex1(2,2)$	Point	100	10,000
$Ex1(2,2)$	Distribution	100	10,000

Table 4: Deriving approximate self-confirming (SC) price predictions in complementary environments: Environments in which the maximum number of iterations or the number of games per iteration deviates from the standard setting (100 and 1 million respectively). Point predictions in rows $U(5,5)$a and $U(5,5)$ differ in the initial prediction (see Table 5 for details). In environment $E(7,9)$, the maximum number of iterations was originally set to 100, but the simulation was aborted by the system after iteration 87.

competitive-equilibrium predictions we have obtained. All these predictions are based on 25,000 draws from the corresponding preference distribution.

1.3.6 Self-Confirming Prediction

The self-confirming (SC) price predictions are predictions that on average are correct if all agents use them. In our empirical-game analysis, we find approximate self-confirming prices by following a simple iterative process until it reaches a fixed point. See Wellman et al. [2008], Section 5.3 for formal definitions and derivation procedure of approximate self-confirming point and distribution predictions. In this section, we present SC predictions that we derived for a number of complementary environments. In substitutable environments, we analyzed a modification of SC prediction that we believe has a higher potential in such environments (see Section 1.5).

In Tables 5 and 6 below, we report approximate SC point predictions and the means of approximate SC distribution predictions. In environments in which the iterative process did not reach a fixed point, the reported price vector is the mean of the last-iteration final-price distribution. In all cases but a few exceptions, the maximum number of iterations is 100 and the number of game instances per iteration is 1 million ($N = 1,000,000$). The exceptions are given in Table 4. Environments $Ex1(2,2)$ and $Ex2(m,n)$ refer to the specific preferences described in Wellman et al. [2008], Example 1 and Example 2 respectively.

1.3.7 Gaussian Distribution Prediction

Let $\boldsymbol{\mu}$ be a vector of expected final prices and $\boldsymbol{\sigma}$ be a vector of standard errors. We can approximate the final-price (marginal) distribution of good $i \in \{1,\dots,m\}$ with a bell-shaped distribution defined on $\{0,\dots,V\}$ and centered around μ_i. We draw $N = 1,000,000$ samples from $N(\mu_i, \sigma_i)$ for each $i \in \{1,\dots,m\}$, round the prices and discard those outside the $[0,\ V]$ interval.[6] Then we compute the final-price probabilities according to Equation (2). We denote the Gaussian prediction by $G(x,y)$, where x and y label the expected-price and the standard-error vectors respectively. For example, the Gaussian prediction of $\mathrm{PP}(G(\mu(EPE_u), \sigma(EPE_u)))$ is based on the means and standard errors of the F^{EPE_u} marginals; the prediction of $\mathrm{PP}(G(\boldsymbol{\pi}^{EDPE_u}, \sigma(EPE_u)))$ is based on vector $\boldsymbol{\pi}^{EDPE_u}$ the standard errors of F^{EPE_u}.

1.3.8 Degenerate Distribution Predictions

Let $\boldsymbol{\pi}$ be an m-element price vector. If the prices are all integers, we create for each good $i \in \{1,\dots,m\}$ an artificial empirical sample in which these prices each occur $N = 1,000,000$ times and the rest of the integer prices in $\{0,\dots,V\}$ never occur. We then compute the marginal PDFs according to Equation (2). If the price of good i in a particular price vector $\boldsymbol{\pi}$ that we want to use is non-integer, we either split a one-million sample between the closest integer prices so that the distribution mean equals the original price $\boldsymbol{\pi}^i$ or simply round the price before generating the distribution. In the former two cases, we denote the degenerate-distribution predictor by $\mathrm{PP}(F(\boldsymbol{\pi}))$. For example, the mean prediction of $\mathrm{PP}(F(\boldsymbol{\pi}^{SB_u}))$ equals exactly $\boldsymbol{\pi}^{SB_u}$. In the latter case, we mark the price vector with a prime. For example, $\mathrm{PP}(F(\boldsymbol{\pi}^{EPE'_u}))$ is based on the rounded $\boldsymbol{\pi}^{EPE_u}$.

[6] One obvious shortcoming of this approach is that the mean of the resulting distribution is different from μ_i, unless $\mu_i = \frac{V}{2}$.

(Complementary) Environment	Final-price distribution mean in the last iteration (= approximate SC point price prediction if a fixed point is reached)
$E(3,3)$	14.32070300 5.43597000 1.65065800
$E(3,5)$	20.097089 8.624289 2.969246
$E(3,8)$	25.343538 11.519247 4.436222
$E(5,3)$	12.329149 4.789413 2.034372 0.920293 0.355297
$E(5,5)$	18.071412 7.883137 3.914884 1.871562 0.670457
$E(5,7)$	22.0834 9.942956 5.353831 2.745027 0.979679
$E(5,8)$	23.794642 10.839511 5.923359 3.117401 1.106578
$E(7,3)$	11.625674 4.438999 1.969657 1.004983 0.570509 0.298123 0.103706
$E(7,6)$	19.742276 8.646408 4.590535 2.670062 1.521491 0.807157 0.322996
$E(7,8)$	23.353738 10.480426 5.90957 3.605955 2.166913 1.181388 0.466163
$E(7,9)$	24.817549 11.222740 6.431461 4.046888 2.487224 1.376034 0.543647
$U(3,3)$	12.988021 6.737967 2.139724
$U(3,5)$	17.741942 9.784769 3.421106
$U(3,8)$	21.786462 11.729432 4.3666
$U(5,3)$	9.532673 6.235117 3.757756 2.012286 0.777861
$U(5,5)$a	12.9295 8.57558 5.33419 2.9881 1.1773
$U(5,5)$	13.038192 8.667978 5.424112 3.034694 1.183108
$U(5,7)$	14.995025 9.559724 5.921804 3.343851 1.324034
$U(5,8)$	15.477431 9.437286 5.634767 3.133098 1.219619
$U(7,3)$	7.764757 5.534822 3.850967 2.640666 1.69385 0.988708 0.443689
$U(7,6)$	11.795494 8.27202 5.799993 4.025271 2.662394 1.566832 0.696669
$U(7,8)$	12.352057 7.917285 5.051705 3.297262 2.108076 1.232187 0.550471
$U(7,9)$	11.876186 6.932236 3.927583 2.338603 1.429579 0.84724 0.384005
$Ex1(2,2)$	14.74893 14.25053
$Ex2(2,2)$	13.919495 13.240667
$Ex2(2,5)$	15.708044 15.046278
$Ex2(5,2)$	6.699457 6.452252 6.08019 5.792035 5.595775
$Ex2(5,5)$	8.060013 7.599247 7.432643 7.302775 7.212703

Table 5: Deriving approximate self-confirming (SC) point price predictions in complementary environments: Distribution means in the last iteration. Price oscillation persisted only in $Ex1(2,2)$. In the rest of the environments, the price vectors satisfy the definition of the approximate SC point price prediction. We used the zero initial prediction PP($\boldsymbol{\pi}^{Zero}$) (see Section 1.3.1) to derive all the predictions in this table except that in row $U(5,5)$. In row $U(5,5)$, the initial prediction is the average of the last 50 (of 70) iterations used to generate the prediction in row $U(5,5)$a.

(Complementary) Environment	KS_{marg} distance in last iteration	Last iteration	Final-price distribution mean in the last iteration (= mean of approximate SC distribution prediction if a fixed point is reached)
$E(3,3)$	-0.001465	100	12.210307 3.804257 1.20171
$E(3,5)$	-0.001836	100	17.862417 6.549359 2.168415
$E(3,8)$	0.000899	100	23.397972 9.344890 3.401826
$E(5,3)$	-0.001255	100	10.517042 3.447020 1.518373 0.764377 0.329298
$E(5,5)$	0.00179	100	16.136437 6.078855 2.952896 1.474867 0.577879
$E(5,7)$	-0.001581	100	20.332879 8.088887 4.194826 2.147838 0.820458
$E(5,8)$	-0.0011	100	22.115973 8.95706 4.752177 2.475092 0.949961
$E(7,3)$	0.0014	100	9.918667 3.202395 1.478728 0.824387 0.502904 0.295518 0.132386
$E(7,6)$	0.001252	100	17.915665 6.844783 3.585964 2.086601 1.238261 0.717492 0.307282
$E(7,8)$	-0.002348	100	21.695038 8.673572 4.751927 2.858787 1.748027 1.016776 0.427326
$E(7,9)$	0.000955	100	23.301749 9.48413 5.281892 3.223573 1.993878 1.153974 0.476095
$U(3,3)$	0.001499	100	10.40263 4.57718 1.517381
$U(3,5)$	0.001845	100	15.238303 7.293387 2.51608
$U(3,8)$	0.005117	100	19.764804 9.706572 3.668295
$U(5,3)$	0.001061	100	7.324368 4.416705 2.687958 1.535882 0.688139
$U(5,5)$	0.007133	6	10.76794 6.545991 4.083502 2.344444 1.026963
$U(5,7)$	-0.00242	100	13.210096 7.948834 5.065203 2.958329 1.290806
$U(5,8)$	0.003864	100	14.22282 8.392991 5.3959 3.188773 1.396048
$U(7,3)$	-0.002343	100	5.957835 4.023661 2.834652 2.001679 1.36823 0.871118 0.432955
$U(7,6)$	0.002896	100	9.741781 6.460071 4.586855 3.258838 2.227502 1.400481 0.690221
$U(7,8)$	-0.002131	100	11.53922 7.439269 5.282462 3.779255 2.603905 1.642525 0.81079
$U(7,9)$	-0.002105	100	12.288118 7.77526 5.529366 3.96556 2.739224 1.727645 0.853398
$Ex1(2,2)$	-0.5055	100	8 7.5024
$Ex2(2,2)$	-0.20994	100	10.826383 10.185083
$Ex2(2,5)$	0.172697	100	12.206948 11.591485
$Ex2(5,2)$	0	2	1 0 0 0 0
$Ex2(5,5)$	0	2	1 0 0 0 0

Table 6: Deriving approximate self-confirming (SC) distribution price predictions in complementary environments: Distribution means in the last iteration. We used uniform initial predictions F^U (see Section 1.3.3) to derive all the distribution predictions in this table except that in row $U(5,5)$, for which we used the SB-based baseline prediction F^{SB_u} (see Section 1.3.4). The KS_{marg} convergence criterion in environment $U(5,5)$ is 0.01, i.e., we programmed the simulation to stop if the KS_{marg} distance between the current- and previous-iteration final-price CDFs for all goods was below 0.01. The criterion was satisfied at iteration 6 in environment $U(5,5)$. For the rest of the environments in this table, the KS_{marg} criterion is 0.00001. This threshold was never reached in any other U- or E-environments, and the simulation stopped at the threshold on the number of iterations, which equals 100 in all the environments. However, KS_{marg} reached 0.01 within the first 11 iterations in all the U-environments and within the first 7 iterations in all the E-environments.

1.4 Demand-Reduction Strategy

The demand-reduction strategy family, DR(κ), was introduced for homogeneous-good environments (see Wellman et al. [2008]). In a homogeneous-good environment, each auction sells one unit of a homogeneous indivisible good, and the bidders' marginal value for one more unit of the good is weakly decreasing. We implemented such preferences by randomly drawing marginal values v_k for the kth unit from $U[0, v_{k-1}]$, where v_0 is a uniform upper bound on the marginal value of a unit, which equals 127 in our empirical-game analysis.

DR(κ)'s perceived price of the unit with the lth lowest *myopically perceived price* (see Section 1.1) is given by

$$\rho_l(\boldsymbol{B}) \equiv \begin{cases} \beta_l + \kappa(l-1) & \text{if winning the unit} \\ \beta_l + 1 + \kappa(l-1) & \text{otherwise,} \end{cases} \tag{3}$$

where $\boldsymbol{\beta}$ is the vector of current bid prices. Thus, the parameter $\kappa \in 0, \ldots, 127$ defines the degree of the agent's demand reduction. An agent with a larger κ bids on fewer units. When $\kappa = 0$, the agent's bidding behavior is equivalent to that of SB (see Section 1.1). In the other extreme case when $\kappa = 127$, the agent never bids at all. For more details on the strategy, see Wellman et al. [2008].

1.5 Own-Effect Price-Prediction Strategy

The own-effect price predictor, OEPP($\boldsymbol{\pi}^x$), is designed for homogeneous-good environments (see Section 1.4; for analysis, see Wellman et al. [2008]) and is a modification of the price predictor (see Section 1.3). Its prediction is an $m \times m$ matrix of *predicted own-effect prices*. Each element of the matrix, which we denote $\pi_{iq}(\boldsymbol{B})$, is a predicted final price of unit i given that the agent tries to win q units and its information state at the current round is $\boldsymbol{B}$. In our analysis, the *initial* price prediction is equal across auctions: $\pi_{iq}(\varnothing) = \pi_{jq}(\varnothing)$ for all i and j for all purchase sizes q. In other words, the initial-prediction matrix consists of m identical rows. We label it by $\boldsymbol{\pi^x}$, where the subscript x labels initial predictions. Of particular interest is what we call *self-confirming own-effect prices*, which we describe in the following section.

1.5.1 Self-Confirming Own-Effect Prices

Self-confirming own-effect prices satisfy the condition that if one of the agents (the "explorer") bids to win q units ignoring its preferences and the other agents "exploit" their own-effect price predictions, that prediction on average *is correct* for all q. See Wellman et al. [2008], Section 7.3 for a formal definition and the derivation procedure. We denote the self-confirming own-effect price matrix by $\boldsymbol{\pi}^{SC}$ and the self-confirming own-effect price-prediction strategy by OEPP($\boldsymbol{\pi}^{SC}$).

We have implemented OEPP($\boldsymbol{\pi}^{SC}$) for a homogeneous-good environment with 5 units and 5 agents (see Wellman et al. [2008]). We set the initial predictions to a matrix of zeros, and simulated 10 thousand games for every purchase size of the explorer. The explorer's purchase size was changed 500 times, i.e., each purchase size was updated 100 times. We created 10 approximate self-confirming own-effect predictions, which we report in Table 7. Since any own-effect initial-prediction matrix consists of m identical rows, we present one row per prediction. In our empirical-game analysis in Wellman et al. [2008], we used the average of these 10 predictions as $\boldsymbol{\pi}^{SC}$.

2 Restricted SAA Games

2.1 53-Strategy Game For 5×5 Uniform Complementary Environment

In Table 8 we describe 53 strategies that we constructed for the 5×5 uniform environment ($U(5,5)$) analyzed in Osepayshvili et al. [2005] and Wellman et al. [2008]. In Table 9 we report all initial point predictions for the price-predicting strategies derived for this environment.

2.2 Strategies For Alternative Complementary Environments

For 11 alternative complementary environments ($E(3,3)$, $E(3,8)$, $E(5,8)$, $U(3,3)$, $U(3,5)$, $U(3,8)$, $U(5,3)$, $U(5,8)$, $U(7,3)$, $U(7,6)$, $U(7,8)$), we have evaluated 27 profiles: one with all PP(F^{SC}), and for each of 26 other strategies s, one profile with $(n-1)$ PP(F^{SC}) and one s, where s is a strategy from Table 10. Also, for 16 alternative complementary environments reported in Table 2 (all environments other than

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 7: Approximate self-confirming own-effect predictions and their average. We used the average as $\boldsymbol{\pi}^{SC}$ in our empirical-game analysis in Wellman et al. [2008]. The columns are the possible target purchase sizes ($q \in \{1, \ldots, m\}$, where $m = 5$ auctions). V=127 in our analysis of homogeneous-good environments.

Strategy Family	Number of Representatives in Game	Strategies
Straightforward Bidder	1	SB
Sunk-aware Agent	20	SA(k) $k = 0, 0.05, 0.1, \ldots 0.95$
Point Price Predictor	13	PP($\boldsymbol{\pi}^{\infty}$), PP($\boldsymbol{\pi}^{SB_u}$), PP($\boldsymbol{\pi}^{SB_u}$) w/ P.O., PP($\boldsymbol{\pi}^{EPE_u}$), PP($\boldsymbol{\pi}^{EPE^*_u}$), PP($\boldsymbol{\pi}^{EPE^{**}_u}$), PP($\boldsymbol{\pi}^{EPE_e}$), PP($\boldsymbol{\pi}^{EPE^*_e}$), PP($\boldsymbol{\pi}^{EDPE_u}$), PP($\boldsymbol{\pi}^{EDPE^*_u}$), PP($\boldsymbol{\pi}^{EDPE_e}$), PP($\boldsymbol{\pi}^{EDPE^*_e}$), PP($\boldsymbol{\pi}^{SC_u}$)
Price Distribution Predictor	19	PP(F^{Zero}),PP(F^{U}) PP(F^{SB_u}),PP(F^{SC_u}), PP(F^{EPE_u}), PP($G(\mu(EPE_u), \sigma(EPE_u))$), PP($G(\mu(SB_u), \sigma(SB_u))$), PP($G(\boldsymbol{\pi}^{EDPE_u}, \sigma(EPE_u))$), PP($G(\boldsymbol{\pi}^{EDPE_u}, \sigma(SB_u))$), PP($G(\boldsymbol{\pi}^{SC_u}, \sigma(EPE_u))$), PP($G(\boldsymbol{\pi}^{SC_u}, \sigma(SB_u))$), PP($F(\boldsymbol{\pi}^{EPE_u})$), PP($F(\boldsymbol{\pi}^{EPE'_u})$), PP($F(\boldsymbol{\pi}^{EDPE_u})$), PP($F(\boldsymbol{\pi}^{EDPE'_u})$), PP($F(\boldsymbol{\pi}^{SB_u})$), PP($F(\boldsymbol{\pi}^{SB'_u})$) PP($F(\boldsymbol{\pi}^{SC_u})$), PP($F(\boldsymbol{\pi}^{SC'_u})$)

Table 8: 53 strategies for the complementary 5×5 uniform-environment game (complementary $U(5,5)$). We report the strategies in column 3 and the strategy family to which they belong in column 1. Column 2 is the total number of strategies from each strategy family. P.O. refers to *participation-only* prediction (see MacKie-Mason et al. [2004]). Different point predictions obtained using the same (non-deterministic) algorithm are marked by asterisks.

Predictions/Good	1	2	3	4	5	Appendix Section
	Initial Point Predictions					
$\boldsymbol{\pi}^{\infty}$	1000	1000	1000	1000	1000	1.3.2
$\boldsymbol{\pi}^{SB_u}$	14.8	10.7	7.6	4.6	1.9	1.3.4
$\boldsymbol{\pi}^{SC_u}$	13.0	8.7	5.4	3.0	1.2	1.3.6
$\boldsymbol{\pi}^{EPE_u}$	16.6	10.8	6.5	3.1	0.7	1.3.5
$\boldsymbol{\pi}^{EPE_u^*}$	16.5	10.7	6.4	3.1	0.8	1.3.5
$\boldsymbol{\pi}^{EPE_u^{**}}$	26.0	14.2	6.9	2.5	0.3	1.3.5
$\boldsymbol{\pi}^{EPE_e}$	6.0	4.1	1.8	0.6	0.1	1.3.5
$\boldsymbol{\pi}^{EPE_e^*}$	30.5	11.9	6.0	2.7	0.4	1.3.5
$\boldsymbol{\pi}^{EDPE_u}$	20.0	12.0	8.0	2.0	0.0	1.3.5
$\boldsymbol{\pi}^{EDPE_u^*}$	20.8	11.4	8.2	1.8	0.0	1.3.5
$\boldsymbol{\pi}^{EDPE_e}$	25.0	10.0	5.1	0.9	0.0	1.3.5
$\boldsymbol{\pi}^{EDPE_e^*}$	24.5	10.5	5.5	1.5	0.0	1.3.5

Table 9: Initial predictions (rounded to one decimal place) for the 5×5 uniform complementary environment (complementary $U(5,5)$). Column 1 is the notation, and column 2 is the point-prediction vectors. The goods are numbered from 1 through 5. The monotonicity of the prices is due to the specifics of the scheduling-game preferences (see MacKie-Mason et al. [2004] and Reeves et al. [2005]). For information about price-predicting strategies, see Section 1.3. The subsections most relevant to particular predictions are given in column 3.

Strategy Family	Number of Representatives in Game	Strategies
Straightforward Bidder	1	SB
Sunk-aware Agent	20	SA(k) $k = \{0, 0.05, 0.1, \ldots 0.95\}$
Point Price Predictors	3	PP($\boldsymbol{\pi}^{\infty}$), PP($\boldsymbol{\pi}^{SB}$), PP($\boldsymbol{\pi}^{SC}$)
Price Distribution Predictors	6	PP(F^U), PP(F^{SB})

Table 10: 26 deviators for 11 alternative environments. We report the strategies in column 3 and the strategy family to which they belong in column 1. Column 2 is the total number of strategies from each strategy family.

complementary $U(5,5)$, which include the 11 above-mentioned environments and 5 additional E models), we evaluated complete 7-cliques. In Table 11 we describe the 7-clique restricted games for each of the alternative environments. We suppress the preference-distribution labels in the tables, because they all match the corresponding environment (e.g., PP(F^{SB}) for $E(5,3)$ refers to PP(F^{SB_e}), and PP(F^{SB}) for $U(5,3)$ refers to PP(F^{SB_u})). See Osepayshvili et al. [2005] and Wellman et al. [2008] for analysis.

2.3 51-Strategy Game for 5×5 Uniform Homogeneous Environment

To analyze the 5×5 homogeneous-good environment (Wellman et al. [2008]), we constructed a 51-strategy restricted game described in Table 12.

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 11: 7-clique restricted games for 16 alternative environments. $E(m,n)$ and $U(m,n)$ refer to the environments with m goods, n agents, and exponential and uniform preference distribution respectively (see Section 1.3 for information about the environments).

Strategy Family	Number of Representatives in Game	Strategies
Straightforward Bidder	1	SB
Sunk-aware Agent	1	SA(k) $k = 0.5$
Price Distribution Predictor	1	PP(F^{SB_u})
Demand-Reduction Agent	47	DR(κ) $\kappa = \{1, 2, \ldots 30, 32, 34,$ $36, 38, 40, 44, 48, 50, 52, 56,$ $60, 70, 80, 90, 100, 110, 120\}$
Own-Effect Price Predictor	1	OEPP($\boldsymbol{\pi}^{SC}$)

Table 12: 51-strategy restricted game for the 5×5 homogeneous-good uniform environment (also referred to as substitutable $U(5,5)$). We report the strategies in column 3 and the strategy family to which they belong in column 1. Column 2 is the total number of strategies from each strategy family.

References

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

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.

Anna Osepayshvili, Michael P. Wellman, Daniel M. Reeves, and Jeffrey K. MacKie-Mason. Self-confirming price prediction for bidding in simultaneous ascending auctions. In *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. In *Fourth ACM Conference on Electronic Commerce*, San Diego, June 2003.

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):67–85, March 2005.

Michael P. Wellman. Methods for empirical game-theoretic analysis (extended abstract). In *Twenty-First National Conference on Artificial Intelligence*, pages 1552–1555, Boston, 2006.

Michael P. Wellman, Anna Osepayshvili, Jeffrey K. MacKie-Mason, and Daniel M. Reeves. Bidding strategies for simultaneous ascending auctions. *The B.E. Journal of Theoretical Economics*, 2008.