Bidding Strategies for Simultaneous Ascending Auctions

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

Daniel M. Reeves
Yahoo Research
45 W 18th St, 6th Floor
New York, NY 10011 USA
dreeves@yahoo-inc.com

2007-12-19 19:33

Abstract

Simultaneous ascending auctions present agents with various strategic problems, depending on preference structure. As long as bids represent non-repudiable offers, submitting non-contingent bids to separate auctions entails an *exposure problem*: bidding to acquire a bundle risks the possibility of obtaining an undesired subset of the goods. With multiple goods (or units of a homogeneous good) bidders also need to account for their own effects on prices. Auction theory does not provide analytic solutions for optimal bidding strategies in the face of these problems. We present a new family of decision-theoretic bidding strategies that use probabilistic predictions of final prices: *self-confirming distribution-prediction* strategies. Bidding based on these is provably not optimal in general. But evidence using empirical game-theoretic methods we developed indicates the strategy is quite effective compared to other known methods when preferences exhibit complementarities. When preferences exhibit substitutability, simpler *demand-reduction* strategies address the own price effect problem more directly and perform better.

1 Introduction

A *simultaneous ascending auction* (SAA) [Cramton, 2005] allocates a set of m related goods among n agents via separate, concurrent English auctions for each good. This is characteristic of a variety of related but not identical real-world auctions, such as bidders participating in concurrent independent auctions on eBay, power markets, spectrum auctions in many countries, and other explicitly designed trading environments

[Milgrom, 2003]. Moreover, some of the key strategic issues presented by SAAs apply whenever there are concurrent markets for interrelated goods, even if those markets are not formal auctions.

Simultaneity is significant only if demands (or supplies) for the various goods are interrelated. We address here some of the challenges bidders with such demands face when formulating their strategies for participation in SAAs. Interrelated demands generally exhibit *complementarity* or *substitutability* (or both), each of which induces characteristic bidding problems.

To study bidding strategies in the face of the strategic challenges presented by complementarity or substitutability, we intentionally abstract from any single application. There are features specific to spectrum auctions, for example, that we do not address, just as there are unaddressed features specific to simultaneous eBay auctions and other particular SAA environments. In hope of producing results generalizable to a range of applications, we analyze a generic SAA exhibiting a few important characteristics that are common across most specific settings.

Complementarity manifests when an agent's value for a good is greater if it also obtains one or more other goods [Lehmann et al., 2006]. For example, an airline passenger may wish to obtain two connecting segments to complete a trip. Goods exhibit complementarity from the perspective of an agent when her valuation for those goods is *superadditive*. Let X, Y, and Z be sets of goods such that $Y \cup Z = X$ and $Y \cap Z = \emptyset$. Given a quasi-linear valuation function, $v\colon 2^{|X|} \to \mathbb{R}$, that assigns value to possible subsets of X, superadditive preference for Y and Z means that $v(X) > v(Y) + v(Z)$. In other words, the combined bundle X is worth more than the sum of its parts. As a special case, if goods in a set are each worthless without the others, they are *perfect complements*. We say that a valuation exhibits complementarities if there are some subsets of goods for which preference is superadditive.

When the inequality is reversed, the valuation is *subadditive*, which occurs for example when goods are substitutes. Flights on the same route by different airlines would typically be considered substitutes, as would flights to two candidate vacation destinations. Technically, goods are substitutes when raising the price of one does not decrease demand for others—that is, for any optimal bundle before the price increase there is an optimal bundle post-increase that includes at least as much demand for all goods that did not increase in price. Substitutability is a strictly stronger condition than subadditivity [Lehmann et al., 2006]. An important extreme case of substitutability is perfect substitutes or *single-unit demand* [Gul and Stacchetti, 1999], where for all $Y \subseteq X$, $v(Y) = \max_{i \in Y} v(\{i\})$. If, in addition, goods are (for this agent) *homogeneous* then they are 1:1 perfect substitutes.

Concurrent auctions with interdependent goods are strategically challenging because agents bid separately in auctions for each item, but willingness-to-pay depends nontrivially on which combination of items the agent ultimately wins. Since bids represent non-repudiable offers, submitting bids to separate auctions entails an *exposure problem*. With complementarities, if an agent bids on a set of items based on her willingness-to-pay for the set, she may pay more than her valuation for the subset she actually wins. With subadditive preferences, an agent bidding based on willingness-to-pay for individual goods risks paying more for a set than it is worth. The SAA mechanism makes it easy for agents to avoid exposure in the case of substitutes. Since a price increase for one good cannot decrease demand for others, the agents can manage their bids to ensure they are never winning more goods than they want at the current prices. With any violation of substitutability, however, a bidder cannot in general obtain a desired package without incurring some exposure risk.

The exposure problem motivates mechanisms that take complementarities directly into account, such as *combinatorial auctions* [Cramton et al., 2005, de Vries and Vohra, 2003], in which the auction mechanism determines optimal packages based on agent bids over bundles. Although such mechanisms may provide an effective solution in many cases, there are often significant barriers to their application [MacKie-Mason and Wellman, 2005]. Indeed, SAA-based auctions are often deliberately adopted, despite awareness of strategic complications [Milgrom, 2003, McAfee and McMillan, 1996].

A second strategic problem for bidders is accounting for *own price effects*: the impact of their own bids on resulting prices. For example, a bidder winning q units may find that bidding less than her incremental value for the $q + 1$st unit results in a price sufficiently lower for the first q units that the inframarginal surplus gain exceeds the incremental surplus from winning the $q + 1$st unit. The strategy of shading bids to exploit this benefit is known as *demand reduction* [Ausubel and Cramton, 2002, Weber, 1997].

Given exposure and own price effects, it is clear that bidding willingness-to-pay is generally not optimal. Worse for designers, researchers, and bidders, auction theory to date [Krishna, 2002] has little to say about how one should bid in simultaneous markets with substitutes or complements. There exists no useful analytical characterization of equilibria for SAA games, and indeed deriving such results appears intractable for nontrivial environments. Moreover, the best-response strategies to even simple specified bidding policies can be surprisingly complex [Reeves et al., 2005]. Simulation studies shed light on some strategic issues [Csirik et al., 2001], as have accounts of strategies employed in specific auctions [Cramton, 1995, Weber, 1997], but the game is too complex to admit definitive strategic recommendations.

We employ a different approach to analyze bidding strategies, which we elsewhere describe as a *computational reasoning* [MacKie-Mason and Wellman, 2005] or *empirical game-theoretic* methodology [Wellman, 2006] for analyzing mechanisms and strategies. We begin with an explicit formulation of the resource allocation problem, generate a set of candidate parametrized strategies, then simulate the game for various profiles of strategy parameters. Through simulation, we in effect convert an extensive-form game of incomplete information with high-dimensional strategy space into a normal-form game over the restricted set of strategies defined by the instances of strategy parameters explored. We then use standard tools to solve the restricted-strategy (yet often still quite large) normal-form games, and analyze the results. For the families of candidate strategies we study, we are able to characterize those which participate in equilibria of the transformed game, and the quality of the resulting outcomes.

One advantage of this fundamentally empirical method is that if others believe they have superior strategies, it is straightforward to apply the method *incrementally* to evaluate the new candidates with respect the best-performing strategies known to date. This is important because the SAA environment is so complex, and in any SAA the specific rules may call for variations on the basic strategy family we study. For example, some auctions impose activity rules, which introduces an eligibility management problem into the design of bidding strategies. Budget constraints may also affect the design of bidding strategies. We do not claim that our present analysis covers the entire space of bidding-strategy design for SAAs. We do claim that our method allows us to accumulate empirical learning on strategy performance in a generic SAA setting, and that strategic lessons from this environment will be a useful starting point for those designing strategies for the particular SAA rules they face.

We next proceed with a formal specification of the problem and of the generic SAA mechanism we study.

2 The Simultaneous Ascending Auctions (SAA) Domain

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 $j \in \{1, \ldots, n\}$,[1] and a specification of the SAA mechanism rules. In general the SAA mechanism comprises m separate auctions, one for each good, 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.

An agent's *current information state*, $\boldsymbol{B}$, comprises the current bid prices, $\boldsymbol{\beta}$, along with a bit vector indicating which goods the agent is currently winning. Let $\boldsymbol{\mathcal{B}}$ denote the set of possible current information states. A *local bidding strategy* is a mapping $\boldsymbol{\mathcal{B}} \to \boldsymbol{b}$, where the bid vector $\boldsymbol{b}$ specifies a bid for each of the m auctions. More generally, an agent's bidding strategy maps the *history* of information states to bids. For the present work, we limit consideration to local bidding strategies. This is a substantive limitation, ruling out, for example, methods that infer other agent's types from dynamic price patterns, or strategies that punish other's behavior. Nevertheless, the strategic issues we consider primary can be addressed at the level of local bidding strategies, and thus we take the simplification achieved through ignoring history to be worthwhile.[2]

Submitting an inadmissible bid (e.g., $b_i = 0$) is equivalent to not bidding. An agent's payoff—also referred to as its *surplus*—is defined by the auction outcomes,

[1] We may include in the type distribution *Nature's type* which determines the random tie-breaking when agents place identical bids.

[2] Assuming that agents submit bids for a subset of goods at the minimum increment, the size of the strategy space is $|\boldsymbol{\mathcal{B}}|2^m$. Conditioning on a history of length t would expand this space by a factor of 2^{mt}.

namely, the set of goods it wins, X, and the final prices, $\boldsymbol{p}$:

$$\sigma(X, \boldsymbol{p}) \equiv v(X) - \sum_{i \in X} p_i. \tag{1}$$

The following sections describe the family of strategies we consider in this study. As explained below, our strategy class is restricted, but covers a wide range of bidding strategies previously explored in the literature.

3 Perceived-Price Bidding Strategies

If an agent knew the final prices of all m goods and if those prices did not depend on its own bidding strategy, then its optimal strategy would be clear: bid on a subset of goods that maximizes its surplus at known prices. When prices are uncertain or bid-dependent, this is not optimal, but may nevertheless serve as a useful starting point. In this section, we define a class of bidding strategies that generalizes this approach by selecting a subset of goods that maximizes surplus at *perceived* prices.

Definition 1 (Perceived-Price Bidder) *A* perceived-price bidder *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,* $\rho(\boldsymbol{B})$*. It 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.[3] *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.*

A perceived-price bidding strategy is defined by how the agent constructs the perceived price from its information state. We now define two versions of the function ρ, corresponding to perceived-price bidding strategies well-studied in prior literature. In Section 4 we define the newer price-prediction perceived-price strategies we analyze in this article. Our discussion focuses on the particularly challenging case of superadditive preference—complementary goods. We return to address the case of substitutable goods in Section 7.

3.1 Straightforward Bidding

One example of a perceived-price bidder is the widely studied *straightforward bidding* (SB) strategy.[4] An SB agent sets $\rho(\boldsymbol{B})$ to *myopically perceived prices*: the bid price

[3] More precisely: when multiple subsets tie for the highest surplus, the agent chooses the smallest. If the smallest subset is not unique it picks the subset whose bit-vector representation is lexicographically greatest. (The bit-vector representation $\boldsymbol{\omega}$ of $X \subseteq \{1, \ldots, m\}$ has $\omega_i = 1$ if $i \in X$ and 0 otherwise. For example, the bit-vector representation of $\{1, 3\} \subseteq \{1, 2, 3\}$ is $\langle 1, 0, 1 \rangle$.) This tie-breaking scheme is somewhat arbitrary, and we expect alternative choices would be inconsequential. We describe our version here in detail to facilitate replication of our experimental results.

[4] We adopt the terminology introduced by Milgrom [2000]. The same concept is also referred to as "myopic best response", "myopically optimal", and "myoptimal" [Kephart et al., 1998].

for goods it was winning in the previous round and the ask price for the others:

$$\rho_i(\boldsymbol{B}) = \begin{cases} \beta_i & \text{if winning good } i \\ \beta_i + 1 & \text{otherwise,} \end{cases} \tag{2}$$

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

Straightforward bidding is a reasonable strategy in some environments. When all agents have single-unit demand, and value every good equally (i.e., the goods are all 1:1 perfect substitutes), 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 [2006] show that straightforward bidding is a perfect Bayes-Nash equilibrium.

If agents have additive utility, i.e., $v(Y) = \sum_{i \in Y} v(\{i\})$, then they can treat the auctions as independent and in this case too, SB is in equilibrium. To see this, consider the case that all other agents are playing SB with additive preference. Then your bid in one auction does not affect your surplus in another. This implies the auctions can be treated independently and SB is a best response.

The degenerate SAA with $m = 1$, i.e., a single ascending auction, is strategically equivalent to a second-price sealed-bid auction [Vickrey, 1961]. In other words, SB is a weakly dominant strategy in a single ascending auction, similarly to "truth-telling" in a second-price sealed-bid auction.[5] For $m > 1$, however, the joint strategy space allows *threats* such as "if you raise the price on my good I will raise it on yours." These will then support demand-reduction equilibria, even in the additive case. Thus, although SB is a good strategy and is in equilibrium for some special-case environments without complementarities, it is not (even weakly) dominant.

Up to a discretization error, the allocation in an SAA with single-unit demand is efficient when agents follow straightforward bidding. It can also be shown [Bertsekas, 1992, Wellman et al., 2001] that the final prices will differ from the minimum unique equilibrium prices by at most $\min(m, n)$ times the bid increment. The value of the allocation, defined to be the sum of the bidder surpluses, will differ from the optimal by at most the bid increment times $\min(m, n)(1 + \min(m, n))$.

Unfortunately, none of these properties hold for general preferences. The final SAA prices 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]. And most importantly, SB need not be a Nash equilibrium.

Example 1 *There are two agents, with values for two goods as shown in Table 1. One admissible straightforward bidding path*[6] *leads to a state in which agent 2 is winning both goods at prices (15,14). Then, in the next round, agent 1 would bid 15 for good 2. The auction would end at this point, with agent 1 receiving good 2 and agent 2 receiving good 1, both at a price of 15.*

[5] Technically, this equivalence applies to a strategically restricted version of the ascending auction which does not allow arbitrary bids above the ask price (and raises the ask price continuously rather than discretely). Otherwise, there exist strategies (albeit pathological) to which SB is not a best response. For example, suppose my policy is to not bid more than \$100 unless the bidding starts lower, in which case I will keep bidding indefinitely. The best response to such a strategy requires *jump bidding*.

[6] The realized progression of the SAA protocol depends on tie-breaking.

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

Table 1: A simple problem illustrating the pitfalls of SB (Example 1).

In this example, SB leads to a result with total allocation value 20, whereas the optimal allocation would produce a value of 30. We can construct slightly more complex examples by adding goods and agents, enabling us to magnify the suboptimality to an arbitrary degree.

We see that straightforward bidding fails to guarantee high quality allocations. It is also easy to show that straightforward bidding is not an equilibrium strategy in general. Consider again Example 1. If agents follow the SB strategy, the mechanism reaches quiescence at prices $\{15, 15\}$. However, it is not rational for agent 2 to stop at this point. If, for example, agent 2 continued bidding, prices would reach $\{21, 20\}$ with agent 2 winning both goods, and the auction would end (assuming agent 1 plays SB). Agent 2 would be better off, with a surplus of -11 rather than -15.

It is clear that SB is not a reasonable candidate for a general strategy in SAA. We show next how a simple parametric generalization to SB can address a key strategic shortfall.

3.2 Sunk-Awareness

We showed in Example 1 that in some problems agents following a straightforward bidding strategy may stop bidding prematurely. We now consider why SB is failing in this situation. In a given round, agents following SB bid on the set of goods that maximizes their surplus at myopically perceived prices (current bid or ask prices). If none of the nonempty subsets of goods appear to yield positive net surplus, the agent chooses the empty set, i.e., it does not bid at all, because the alternative is to earn negative surplus. However, this behavior ignores outstanding commitments: the agent may already be winning one or more goods. If the agent drops out of the bidding, and others do not bid away the goods the agent already is winning, then its alternative surplus could be much worse than if it continued to bid despite preferring the empty bundle at current prices. In the case of an agent dropping out of the bidding on some goods in a bundle of perfect complements, its surplus is negative the sum of the bid prices for the goods in the bundle it gets stuck with. This failure of straightforward bidding is due to ignoring the true opportunity cost of not bidding.

We refer to this property of straightforward bidding as "sunk-unawareness" [Reeves et al., 2005]. SB agents bid as if the incremental cost for goods they are currently winning is the full price, β_i. However, if the probability that someone else will outbid the agent for this good is α, then the agent is already committed to an expected payment of $(1-\alpha)\beta_i$. This represents a sunk cost that should not affect rational continuation bidding. We can think of the difference, $\alpha\beta_i$, as a rough measure of the incremental cost the agent incurs by deciding to stay with this good.

To address this limitation of straightforward bidding, we parametrize a family of perceived-price bidding strategies (Definition 1) that permits agents to account to a greater or lesser extent for the true incremental cost of goods they are currently winning. We call this strategy "sunk aware". A sunk-aware agent bids as if the incremental cost for goods it is currently winning is somewhere on the interval of zero and the current bid price.

Our sunk-aware strategies generalize SB's method for choosing the perceived-price vector (Equation 2) through the parameter $k \in [0, 1]$:

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

Using this perceived-price vector to define sunk-aware bidders, Definition 1 above gives us a complete specification of the agent's bidding strategy. 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. Intermediate values are akin to bidding as if the agent puts an intermediate probability on the likelihood of retaining the goods it is currently winning. We treat as a special case agents with single-unit demand: their sunk-aware strategy is to bid straightforwardly ($k = 1$) since for such agents SB is a no-regret strategy.

The sunk-awareness parameter provides a heuristic for a complex tradeoff: the agent's bidding behavior changes after it finds itself exposed to the underlying problem (owning goods for which the agent has lower value if not part of a larger package). In our previous study we experimentally determined good settings of the sunk-awareness parameter in various environments [Reeves et al., 2005].

4 Prediction-Based Perceived-Price Bidding Strategies for SAA

Whenever an agent has non-substitutes preference and chooses to bid on a bundle of size greater than one, it may face exposure. Exposure in SAA is 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 decision-theoretic approach would account for these expected costs and benefits, bidding when the benefits prevail, and cutting losses in the alternative.

Re-consider agent 2's plight in Example 1: following SB it is caught by the exposure problem, stuck with a useless good and a surplus of -15. (Other tie-breaking choices result in different outcomes but all of them leave agent 2 exposed and with negative surplus.) If the agent instead plays a fully sunk-aware strategy the result could be an outcome in which it purchases both goods at prices $\{21, 20\}$ for a net surplus of $30 - 41 = -11$. This is better than using SB, but the agent would fare better still by not bidding at all.

The effectiveness of a particular strategy will in general be highly dependent on the characteristics of other agents in the environment. This observation motivates the use of *price prediction*. We would prefer strategies that employ type-distribution beliefs to guide bidding behavior, rather than relying only on current price information as in the sunk-aware strategies (including SB). Forming price predictions for the goods in SAA is a natural use for type-distribution beliefs. In Example 1, suppose agent 2 could predict with certainty before the auctions start that the prices would total at least 30. 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.

We now propose to improve on SB and sunk-awareness by using explicit price predictions for perceived prices. Let $F \equiv F(\boldsymbol{B})$ denote a joint cumulative distribution function over final prices, representing the agent's belief given its current information state $\boldsymbol{B}$. We assume that prices are bounded above by a known constant, V. Thus, F associates probabilities with price vectors in $\{1, \ldots, V\}^m$.

We next consider two ways to use prediction information to generate perceived prices. We first define a *point prediction* for perceived prices, $\boldsymbol{\pi}$, that anticipates possible exposure risks. Then we define a *distribution prediction* for perceived prices, $\boldsymbol{\Delta}$, that explicitly models uncertainty about the exposure prospects. The distribution approach also supports evaluation of the likelihood that the agent's current winning bids are sunk costs. (As with sunk-awareness, price-prediction strategies for agents with single-unit demand ignore the predictions and play SB.)

Before we define our price-prediction strategies we want to make two points. First, we are not (initially) concerned with *how* the agent formulates her beliefs (price predictions), nor the optimality of the prediction method. Rather, we propose strategies that use *some* beliefs. In our experiments we investigate several different predictors.[7] Second, since these are strategies for bidding in iterative auctions, we face the question of how to update the initial price predictions (based on information revealed by the bidding, for example), we are not concerned in this paper with discovering the optimal updating procedure; again, we are interested in designing strategies that are defined up to the use of *some* belief updating procedure. Indeed, while we experiment with different initial price predictors, in this paper we use only one specific, simple, updating procedure.

4.1 Point Price Prediction

Suppose the agent has (at least) point beliefs about the final prices that will be realized for each good. Let $\boldsymbol{\pi}(\boldsymbol{B})$ be a vector of predicted final prices. Before the auctions begin the price prediction is $\boldsymbol{\pi}(\varnothing)$, where $\varnothing$ is the null information state available pre-auctions.

The auctions in SAA reveal the bid prices each round. Since the auctions are ascending, once the current bid price for good i reaches β_i, there is zero probability that the final price p_i will be less than β_i. We define a simple updating rule using this fact: the current price prediction for good i is the maximum of the initial prediction and the

[7] We believe that finding the *optimal* predictor to use in a particular strategy is likely to be as computationally infeasible as the problem of finding an optimal strategy.

myopically perceived price:

$$\pi_i(\boldsymbol{B}) \equiv \begin{cases} \max(\pi_i(\varnothing), \beta_i) & \text{if winning good } i \\ \max(\pi_i(\varnothing), \beta_i + 1) & \text{otherwise.} \end{cases} \tag{3}$$

Armed with these predictions, the agent plays the perceived-price bidding strategy (Definition 1) with $\rho(\boldsymbol{B}) \equiv \boldsymbol{\pi}(\boldsymbol{B})$. We denote a specific point price-prediction strategy in this family by $\mathrm{PP}(\boldsymbol{\pi}^x)$, where x labels particular initial prediction vectors, $\boldsymbol{\pi}(\varnothing)$. Note that straightforward bidding is the special case of price prediction with the predictions all equal to zero: $\mathrm{SB} = \mathrm{PP}(\boldsymbol{0})$. If the agent underestimates the final prices, it will behave identically to SB after the prices exceed the prediction. If the agent overestimates the final prices, it may stop bidding prematurely.

4.2 Distribution Price Prediction

We generalize the class of price-prediction strategies by taking into account the entire distribution F, rather than just a nominal point estimate (e.g., the expectation of F). We assume the agent generates $F(\varnothing)$, an initial, pre-auction belief about the distribution of final prices.

As with the point predictor, we restrict the updating in our distribution predictor to conditioning the distribution on the fact that prices are bounded below by $\boldsymbol{\beta}$. Let $\Pr(\boldsymbol{p} \mid \boldsymbol{B})$ be the probability, according to F, that the final price vector will be $\boldsymbol{p}$, conditioned on the information revealed by the auction, $\boldsymbol{B}$. Then, with $\Pr(\boldsymbol{p} \mid \varnothing)$ as the pre-auction initial prediction, we define:

$$\Pr(\boldsymbol{p} \mid \boldsymbol{B}) \equiv \begin{cases} \dfrac{\Pr(\boldsymbol{p} \mid \varnothing)}{\sum_{\boldsymbol{q} \geq \boldsymbol{\beta}} \Pr(\boldsymbol{q} \mid \varnothing)} & \text{if } \boldsymbol{p} \geq \boldsymbol{\beta} \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

(By $\boldsymbol{x} \geq \boldsymbol{y}$ we mean $x_i \geq y_i$ for all i.) For (4) to be well defined for all possible $\boldsymbol{\beta}$ we define the price upper bounds such that $\Pr(V, \ldots, V \mid \varnothing) > 0$.

We now use the distribution information to implement a further enhancement to take sunk costs into account in a more decision-theoretic way than the sunk-aware agent. If an agent is currently not winning a good and bids on it, then the expected incremental cost of winning the good is the expected final price, with the expectation calculated with respect to the distribution F. If the agent is currently winning a good, however, then the expected incremental cost of winning that good depends on the likelihood that the current bid price will be increased by another agent, so that the first agent has to bid again to obtain the good. If, to the contrary, it keeps the good at the current bid, the full price is sunk (already committed) and thus should not affect incremental bidding. Based on this logic we define $\Delta_i(\boldsymbol{B})$, the expected *incremental* price for good i.

First, for simplicity, we use only the information contained in the vector of marginal distributions, $(F_1, \ldots, F_m)$, as if the final prices were independent across goods. Define the expected final price conditional on the most recent vector of bid prices, $\boldsymbol{\beta}$:

$$E_F(p_i \mid \boldsymbol{\beta}) = \sum_{q_i=0}^{V} \Pr(q_i \mid \beta_i) q_i = \sum_{q_i=\beta_i}^{V} \Pr(q_i \mid \beta_i) q_i.$$

The expected incremental price depends on whether the agent is currently winning good i. If not, then the lowest final price at which it could win is $\beta_i + 1$, and the expected incremental price is simply the expected price conditional on $p_i \geq \beta_i + 1$,

$$\Delta_i^{\mathrm{L}}(\boldsymbol{B}) \equiv E_F(p_i \mid p_i > \beta_i + 1) = \sum_{q_i=\beta_i+1}^{V} \Pr(q_i \mid p_i > \beta_i + 1) q_i.$$

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

$$\Delta_i^{\mathrm{W}}(\boldsymbol{B}) = (1 - \Pr(\beta_i \mid \beta_i)) \sum_{q_i=\beta_i+2}^{V} \Pr(q_i \mid \beta_i + 2) q_i.$$

The vector of expected incremental prices is then defined by

$$\Delta_i(\boldsymbol{B}) = \begin{cases} \Delta_i^{\mathrm{W}}(\boldsymbol{B}) & \text{if winning good } i \\ \Delta_i^{\mathrm{L}}(\boldsymbol{B}) & \text{otherwise.} \end{cases}$$

The agent then plays the perceived-price bidding strategy (Definition 1) with $\rho(\boldsymbol{B}) \equiv \boldsymbol{\Delta}(\boldsymbol{B})$. We denote the strategy of bidding based on a particular distribution prediction by $\mathrm{PP}(F^x)$, where x labels various pre-auction distribution predictions, $F(\varnothing)$.

5 Some Methods for Predicting Prices in SAA

Definition 1 parametrizes the class of perceived-price bidding strategies, and in Section 4 we define point price and distribution-based prediction methods that construct perceived prices from the agent's information state. The point and distribution predictors are classes of strategies parametrized by the choice of initial prediction—a vector of predicted final prices in the case of the point predictor, or, more generally, a distribution of final prices for the distribution predictor. We now present several ways to obtain an initial prediction. These methods all take as input the problem's type distribution, and so (unlike a particular sunk-awareness setting, for example) are potentially appropriate to apply across different environments.

5.1 Walrasian Equilibrium for Point and Distribution Prediction

The first method we describe predicts that prices will attain a *Walrasian price equilibrium* with respect to the m goods and agent valuation functions over those goods. In practice, to use this method, agents (who know only their own valuation function) will need to have probabilistic beliefs over the valuation functions for other agents.

Given a distribution of agent types we can generalize the price-equilibrium calculation in two ways to allow for probabilistic knowledge of the aggregate demand function. The first is to find the *expected price equilibrium* (EPE): the expectation (over the type distribution) of the Walrasian price-equilibrium vector. The most straightforward way to estimate this is Monte Carlo simulation, sampling from the type distribution. A particular sampled type determines the demand function $\boldsymbol{x}$, which we can then employ in a tatonnement protocol. Repeated sampling of types and application of tatonnement yields a crude Monte Carlo estimate of the expected price equilibrium.

An alternative (which may sometimes be preferred for computational reasons) to estimating a price equilibrium in the face of probabilistic demand is the *expected-demand price equilibrium* (EDPE): the Walrasian price equilibrium with respect to expected aggregate demand. In other words, we calculate or estimate the expected demand function and then apply tatonnement once to find an equilibrium and equilibrium prices as if realized demand were in fact equal to expected demand. We calculate expected demand analytically when possible; otherwise, we can estimate it by Monte Carlo simulation, again sampling from the type distribution.

Either of these generalized Walrasian price-equilibrium methods can be applied to generate point predictions. We denote the expected price-equilibrium point predictor by $\mathrm{PP}(\boldsymbol{\pi}^{\mathrm{EPE}})$ and the expected-demand price-equilibrium point predictor by $\mathrm{PP}(\boldsymbol{\pi}^{\mathrm{EDPE}})$. The method of expected price equilibrium can also be straightforwardly generalized—by tracking the empirical distribution of price equilibria instead of just average prices—to the case of distribution predictors. This predictor is denoted $\mathrm{PP}(F^{\mathrm{EPE}})$.[8]

5.2 Predictions from Historical Data

One natural method for generating an initial prediction is to observe a history of simulated games, and calculate from the outcomes an empirical price distribution for a given strategy profile. Below, we run our game simulator repeatedly, tracking not payoffs but final prices, similar to the tatonnement-sampling approach to estimating equilibrium price distributions. For a point price prediction we compute average final prices, and for a distribution-based prediction we compute final price histograms. We then specify a particular price-prediction strategy by the type of predictor (point vs. distribution) and by a strategy profile from which we glean a distribution of final prices.

As a noteworthy special case of the above, our baseline prediction is the distribution of final prices resulting when all agents follow the SB strategy. We denote the baseline point predictor $\mathrm{PP}(\boldsymbol{\pi}^{\mathrm{SB}})$ and the baseline distribution predictor $\mathrm{PP}(F^{\mathrm{SB}})$.

[8] Unlike the EPE method, which produces a price vector for each sample from the type distribution, the EDPE-method price data is always a price vector, because tatonnement is applied only once at the last step. Therefore, we did not construct *distribution* price-prediction strategies based on the latter; we do construct $\mathrm{PP}(\boldsymbol{\pi}^{\mathrm{EDPE}})$, the EDPE *point* prediction strategies.

5.3 Self-Confirming Price Predictions

Related to predictions from historical data, but somewhat forward-looking, is a price estimate for a prediction strategy based on simulations of itself. We refer to these as *self-confirming* predictions. We begin with the simpler case of point predictions.

Definition 2 (Self-Confirming Point Price Prediction) *Let Γ be an instance of an SAA game. The prediction $\boldsymbol{\pi}$ is a* self-confirming prediction for Γ *iff $\boldsymbol{\pi}$ is equal to the expectation (over the type distribution) of the final prices when all agents play* $\mathrm{PP}(\boldsymbol{\pi})$.

In other words, if all agents use a point price-prediction strategy, then the self-confirming predictions are those that on average *are correct* at the end of the auction.[9] We denote the self-confirming prediction vector by $\boldsymbol{\pi}^{\mathrm{SC}}$ and the self-confirming point prediction strategy by $\mathrm{PP}(\boldsymbol{\pi}^{\mathrm{SC}})$.

The key feature of self-confirming predictions is that agents make decisions based on predictions that turn out to be correct with respect to the type distribution and the assumption that all agents play this particular prediction strategy.[10] Since agents are employing these predictions strategically, we might reasonably expect the strategy to perform well in an environment where its predictions are confirmed.

We next define the concept of a self-confirming *distribution* of final prices in SAA.

Definition 3 (Self-Confirming Price Distribution) *Let Γ be an instance of an SAA game. The prediction F is a* self-confirming price distribution for Γ *iff F is the distribution of prices resulting when all agents play bidding strategy* $\mathrm{PP}(F)$.

The actual joint distribution will in general have dependencies across prices for different goods. We are also 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 4 (Self-Confirming Marginal Distribution) *Let Γ be an instance of an SAA game. The prediction $F = (F_1, \ldots, F_m)$ is a vector of* self-confirming marginal price distributions for Γ *iff for all i, F_i is the marginal distribution of prices for good i resulting when all agents play bidding strategy $\mathrm{PP}(F)$ in Γ.*

5.3.1 Existence of Self-Confirming Predictions

We demonstrate in Section 5.3.2 that we can often find approximately self-confirming point and distribution predictions. However, we first observe that they do not always exist. Consider once again the $m = n = 2$ configuration of Example 1 (Table 1). Versions of this example, in which one agent views the goods as complements and the other as substitutes, are commonly employed to illustrate the absence of a competitive

[9] As described above, our price-prediction strategies perform simple updating based on price-quote information as the auction proceeds. Our self-confirmation notion, however, applies only to initial predictions and final prices–we do not insist that the intermediate updated predictions are also confirmed.

[10] An equilibrium with this feature is sometimes called a "fulfilled expectations equilibrium" [Novshek and Sonnenschein, 1982].

equilibrium [Cramton, 2005, McAfee and McMillan, 1996]. There exist no prices for goods 1 and 2 such that both agents optimize their demands at the specified prices and the markets clear.

Proposition 1 *There exist SAA games for which no self-confirming point price prediction exists, nor do any self-confirming or marginally self-confirming price distributions.*

Proof. Define an SAA game corresponding to the configuration of Table 1. Given a deterministic SAA mechanism (one without asynchrony or random tie-breaking), 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 $\mathrm{PP}(F)$ 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 price distributions that are self-confirming to a reasonable degree of approximation exist for a large class of nondegenerate preference distributions, and can be computed given a specification of the preference distribution. We now present a procedure for deriving such distributions, and some evidence for its effectiveness.

5.3.2 Deriving Self-Confirming Price Predictions

To find approximate self-confirming point predictions, we follow a simple iterative procedure. First, we initialize the predicting agents with some prediction vector (e.g., all zero) and simulate many game instances with the all-predict profile. When average prices obtained by these agents are determined, we replace the initial prediction vector with the average prices and repeat. When this process reaches a fixed point, we have the self-confirming prediction, $\boldsymbol{\pi}^{\mathrm{SC}}$. In Figure 1 we show the convergence to a self-confirming price-prediction vector for a particular SAA game with five agents and five goods. Within 30 iterations the prices have essentially converged, although there is some persistent oscillation. We found that by resetting the vector of predicted prices to equal the averages around which the prices are oscillating, the process converges immediately to a more precise fixed point, which we use as $\boldsymbol{\pi}^{\mathrm{SC}}$.

A similar approach can be applied to derive distribution predictions. Starting from an arbitrary prediction F^0, we run many SAA game instances (sampling from the given preference distributions) with all agents playing strategy $\mathrm{PP}(F^0)$.[11] We record the resulting prices from each instance, and designate the sample distribution observed by F^1. We repeat the process using the new distribution F^t for iteration $t+1$ 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.

[11] In our experiments the initial prediction is zero prices, but our results do not appear sensitive to this.

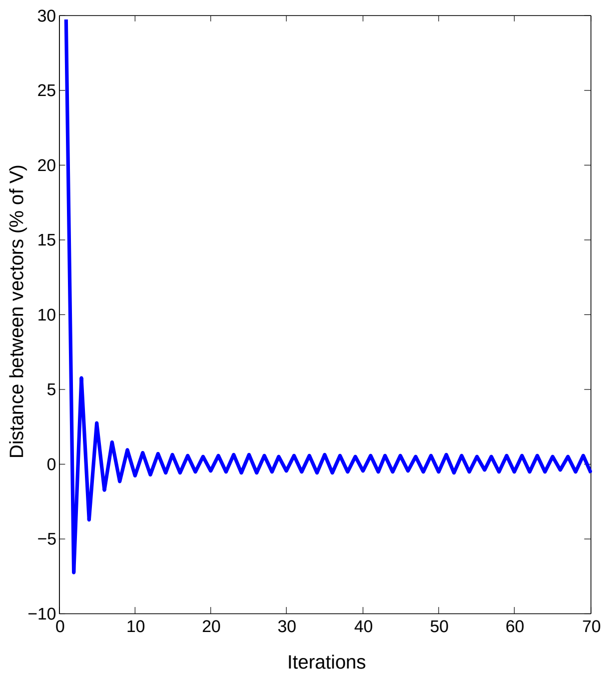

Figure 1: Convergence to a self-confirming price-prediction vector, starting with an initial prediction that all prices would be zero. The prices at each iteration are determined by 500 thousand simulated games. The graph plots the distance between the price vectors in consecutive iterations. We define vector distance as the maximum over pointwise distances, measured as a percentage of the upper bound on the value, V, of a single good. The bound V equals 50 in all of our SAA games with complementary preferences.

We employ the Kolmogorov-Smirnov (KS) statistic as one reasonable measure of similarity of probability distributions, defined as the maximal distance between any two corresponding points in the CDFs:

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

For self-confirming marginal distributions, we take the maximum of the KS distances measured separately for each good: $KS_{\text{marg}} = \max_i KS(F_i, F'_i)$.

Our procedure requires (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, and (4) a smoothing parameter k designating a number of iterations to average over when the procedure reaches the maximum iterations without meeting the threshold. The bound on the number of iterations ensures the 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. We do not, of course, expect the bidding strategy to perform as well when we cannot find a convergent self-confirming distribution and the underlying oscillations are large.

For our empirical analyses, we specify an SAA game based on a scheduling problem in which there are m units (called *time slots*) of a single schedulable resource, indexed $1, \ldots, m$. Each of n agents has a single job that can be accomplished using the resource. Agent j's job requires λ_j time slots to complete, and by accomplishing this job it obtains some value depending on the time it completes. Specifically, if j acquires λ_j time slots by deadline t, it accrues value $v_j(t)$. Deadline values are nonincreasing: $t < t'$ implies $v_j(t) \geq v_j(t')$.

To illustrate, we consider such a scheduling problem with five agents competing for five time slots. We draw job lengths randomly from $U[1, 5]$. We choose deadline values randomly from $U[1, 50]$ then prune to impose monotonicity [Reeves et al., 2005]. We set the algorithm parameters at one million games per iteration, and a 0.01 KS convergence criterion. The predicted and empirical distributions quickly converge, with a KS statistic of 0.007 after only six iterations.

To see if our method produces useful results with some regularity, we applied it to 22 additional instances of the scheduling problem, varying the numbers of agents and goods, and the preference distributions. 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 21 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 constructed 10 instances of the uniform model, comprising various combination of $3 \leq n \leq 9$ and $3 \leq m \leq 7$. In each case, our procedure found 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. We plot the distribution of KS values from these 21 instances in Figure 2.

The 22nd instance was designed to be more challenging: we used the $n = m = 2$ example with fixed preferences described in Table 1. Since there exists no SC distribu-

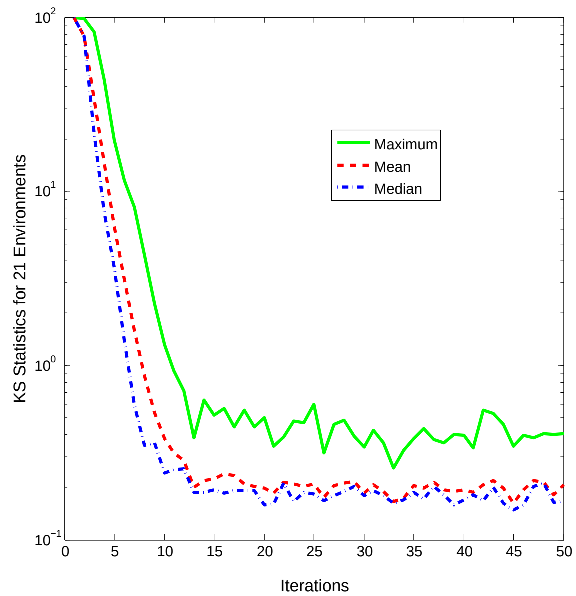

Figure 2: Convergence of iterative SC price estimation.

tion, 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 $k = 10$.

6 Empirical-Game Analysis: Complementary Preferences

We now analyze the performance of self-confirming price distribution predictors in a variety of SAA games, against a variety of other strategies. We use Monte Carlo simulation to estimate the payoff function for an *empirical game*, which maps profiles of agent strategies to expected payoffs for each agent. This approach converts a game in extensive form to normal form in the expected payoffs. We then analyze equilibria in these normal forms. Our methods extend the approach developed in our prior work

[MacKie-Mason et al., 2004, Reeves et al., 2005, Wellman, 2006], and build on ideas from other recent studies in a similar empirical vein [Armantier et al., 2000, Kephart et al., 1998, Walsh et al., 2002]. We emphasize here that all of the analysis below applies directly to the *estimated* empirical game. These correspond to statistical claims about the actual restricted-strategy game, and lead to arguments generalizing the observations to related games.

6.1 Environments and Strategy Space

We studied SAAs applied to market-based scheduling problems, as described in Section 5.3.2. Particular environments are defined by specifying the number m of goods, the number n of agents, and a preference model comprising probability distributions over job lengths and deadline values. The bulk of our computational effort went into an extensive analysis of one particular environment, the $m = n = 5$ uniform model presented above. As described in Section 6.2, the empirical game for this setting provides much evidence supporting the unique strategic stability of $\mathrm{PP}(F^{SC})$. 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 $3 \leq n \leq 8$; goods in $3 \leq m \leq 7$.

To varying degrees, we have analyzed the interacting performance of 53 different strategies. These were drawn from four strategy families described above: SB, sunk-aware, point predictor, and distribution predictor. For each family we varied a defining parameter to generate the different specific strategies.[12] The choice of strategies was based on prior experience. We believe that the set includes the best strategy candidates from the prior literature, though we make no claim to have covered all reasonable variations. Naturally, our emphasis is on evaluating the performance of $\mathrm{PP}(F^{SC})$ in combination with the other strategies.

Given n agents and S possible strategies, the corresponding symmetric normal-form game comprises $\binom{n+S-1}{n}$ distinct strategy profiles. The game size thus grows exponentially in n and S; for the $n = 5$, $S = 53$ game we estimate below, there are over four million different strategy profiles to evaluate. We first illustrate the process for a simpler game, with five agents, each choosing between SB or the baseline point price-prediction strategy $\mathrm{PP}(\boldsymbol{\pi}^{\mathrm{SB}})$ (abbreviated PP). There are six possible profiles which can be described as profiles with j agents playing PP (and the rest SB) for $j = 0, \ldots, 5$. We simulate a large number of games for each profile and average the payoffs for a player of each type (PP, SB). We present the resulting empirical game in Figure 3. For this simple game, we can solve the normal form for a unique pure-strategy Nash equilibrium by inspection, illustrated by the arrows. If all five players choose SB, any one can get a higher expected payoff by deviating to PP. If only one plays PP, a second can beneficially deviate to PP. Likewise for each profile except all playing PP, from which none can gain by deviating to SB, establishing a unique Nash equilibrium.

[12] Space considerations preclude a full description of the 53 strategies here. An appendix with specification of all parameters, including description of the prediction methods used for point and distribution predictors,

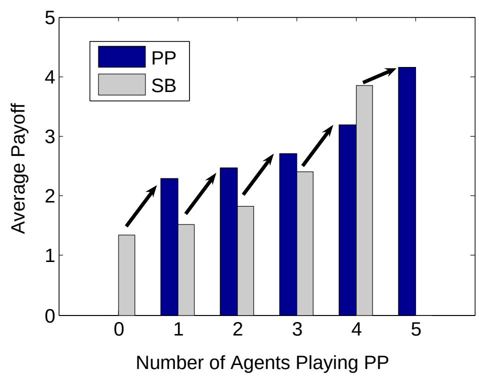

Figure 3: Normal-form payoffs for a 5-player game with 2 strategies. The arrows indicate best responses. All-PP is the unique Nash equilibrium.

6.2 5×5 Uniform Environment

By far the largest empirical SAA game we have constructed is for the SAA scheduling environment discussed in Section 5.3.2, with five agents, five goods, and uniform distributions over job lengths and deadline values. We estimate payoffs empirically for each profile by running millions of simulations of the auction protocol, so estimating the entire payoff function for over 4.2 million strategy profiles is infeasible. However, we can estimate the payoff matrix for subsets of all profiles, and as we describe below, with well-chosen subsets we can reach useful conclusions about equilibria in the 53-strategy game.

Our results are based on estimated payoffs for 4958 strategy profiles, calculated from an average of 10 million samples per profile (with some profiles simulated for as few as 200 thousand games, and some for as many as 200 million, depending on sampling variances). 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-prediction strategy, $\mathrm{PP}(F^{SC})$, would perform well. We have directly verified this: *the profile where all five agents play a pure* $\mathrm{PP}(F^{SC})$ *strategy is a Nash equilibrium of the empirical game*. 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 a strategy s) for n players and S strategies, one needs only S profiles: one for each strategy playing against $n-1$ copies of s. Similarly, to refute the possibility of a particular profile being in Nash equilibrium, we need to find only one profitable deviation profile (i.e., obtained by changing the strategy of one player to earn a higher payoff given the others' strategies).

The fact that $\mathrm{PP}(F^{SC})$ 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.[13] Each clique defines a subgame, for which we have complete payoff information. Within our 4958 profiles we have eight maximal cliques that include strategy $\mathrm{PP}(F^{SC})$. For each of these subgames, $\mathrm{PP}(F^{SC})$ is the only strategy that survives iterated elimination of (strictly) dominated strategies. It follows that $\mathrm{PP}(F^{SC})$ is the *unique* (pure- or mixed-strategy) Nash equilibrium in each of these clique games. 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 $\mathrm{PP}(F^{SC})$ equilibrium.

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 $\mathrm{PP}(F^{SC})$, we have all profile combinations for 49. Based

is available at `http://www-personal.umich.edu/~annaose/papers/saa-appendix.pdf`.

[13]Thus we have a 2-strategy clique if we have estimated all six profiles that five agents can form from these two strategies.

on profiles estimated, we have determined that for any symmetric profile defined by a mixture of one of these pairs, an agent can improve its payoff by a minimum of 0.32 through deviating to some other pure strategy. For reference, the average payoff for the all-PP(F^{SC}) profile is 4.51, so this represents a nontrivial difference.

That is, none of the two-strategy mixtures for which we have data comes very close to equilibrium, further strengthening our confidence in $\mathrm{PP}(F^{SC})$.

Finally, for each of the 4958 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 $\mathrm{PP}(F^{SC})$: $\epsilon = 0$ (confirmed Nash equilibrium of the empirical game);
2. one $\mathrm{PP}(F^{SB})$, four $\mathrm{PP}(F^{SC})$: $\epsilon > 0.13$;
3. two $\mathrm{PP}(F^{SB})$, three $\mathrm{PP}(F^{SC})$: $\epsilon > 0.19$.

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

Our conclusion from these observations is that $\mathrm{PP}(F^{SC})$ 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 analysis of only one particular distribution of preferences, so we now consider other environments.

6.3 Self-Confirming Prediction in Other Environments

To test whether the strong performance of $\mathrm{PP}(F^{SC})$ generalizes across other SAA games, we undertook smaller versions of this analysis on variations of the model above. We explored 17 additional instances of the market-based scheduling problem: eight with the uniform (U), eight with the exponential (E) preference models (3–8 agents, 3–7 goods), and one with fixed preferences, corresponding to the counterexample model of Table 1. For each we derived self-confirming price distributions (failing in the last case, of course), as reported in Section 5.3.2. We also derived price vectors and distributions for the other prediction-based strategies. For 11 of the symmetric games (eight U and three E models), we evaluated 27 profiles: one with all $\mathrm{PP}(F^{SC})$, and for each of 26 other strategies s, one profile with $n-1$ $\mathrm{PP}(F^{SC})$ and one s. For the non-symmetric game with fixed preferences, we evaluated all 53 profiles with at least one agent playing $\mathrm{PP}(F^{SC})$. We ran between two and ten million games per profile in all of these environments.

In eight of the eleven symmetric games, $\mathrm{PP}(F^{SC})$ and $\mathrm{PP}(F^{SB})$ were among top three unilateral deviations from $\mathrm{PP}(F^{SC})$ in the all-PP(F^{SC}) profile. For each of the eleven games, we identified five (additional) top-ranking deviations from $\mathrm{PP}(F^{SC})$ and evaluated complete 7-cliques involving these five strategies, $\mathrm{PP}(F^{SC})$ and $\mathrm{PP}(F^{SB})$ in the respective environments (at least 340,000 samples per profile). We introduced five additional E models, and evaluated all profiles over seven strategies for each of these as well.[14]

[14]For these additional five E models we did not incur the additional computational cost of evaluating all 27 profiles to select best deviations from $\mathrm{PP}(F^{SC})$, which is a somewhat arbitrary procedure for selecting strategies for a clique in any case. For these additional models, we selected the seven candidate strategies based on regularities in the results from the other 11 games described above.

Our results for U and E models are summarized in Table 2. For each case, we report the ϵ that, for the estimated payoff matrix, renders all-PP(F^{SC}) an ϵ-Nash equilibrium. The next two columns report sensitivity information about this figure, given its basis in payoffs estimated from samples. First, since our payoff matrix is estimated (and thus each payoff has a sampling variance), we calculate the expected value $\bar{\epsilon}$ of ϵ with respect to the empirical distributions of the estimated payoffs (assuming that the errors in our payoff estimates are independent, and using the sample variances as population variances). Thus, for example, the environment $E(3, 5)$ has a pure Nash equilibrium of all-PP(F^{SC}) for the estimated payoff matrix, but taking into account sampling variation, on average that profile has an ϵ of 0.005.

Under the same independence assumption, "$\Pr(\epsilon = 0)$" represents the probability that all-PP(F^{SC}) is actually an equilibrium. Finally, for each empirical game with $n \leq 6$ we also obtained a symmetric mixed-strategy Nash equilibrium using replicator dynamics.[15] The rightmost column reports the probability of playing $\mathrm{PP}(F^{SC})$ in the resulting mixture, to evaluate its significance when it does not constitute a pure-strategy equilibrium.

Env(m, n)	ϵ-gain from one-player deviation	$\bar{\epsilon}$-gain adjusted for sampling error	$\Pr(\epsilon = 0)$: Probability of exact Nash equilibrium	Probability of play in rep. dyn. solution
$E(3, 3)$	0	0	1.00	1.00
$E(3, 5)$	0	.005	.600	.996
$E(3, 8)$	.031	.032	0	—
$E(5, 3)$	0	0	1.00	.999
$E(5, 5)$	0	.001	.900	.998
$E(5, 8)$	.029	.031	0	—
$E(7, 3)$	0	.007	.667	.992
$E(7, 6)$	.003	.007	.567	.549
$U(3, 3)$	.097	.099	0	.725
$U(3, 5)$	0	0	1.00	1.00
$U(3, 8)$	.017	.016	0	—
$U(5, 3)$	.103	.103	0	.809
$U(5, 8)$	.047	.048	0	—
$U(7, 3)$	.058	.060	0	.942
$U(7, 6)$	.018	.018	0	.929
$U(7, 8)$	.133	.132	0	—

Table 2: Evaluations of all-PP(F^{SC}) profile for U and E models.

[15] By replicator dynamics we mean an iterative (evolutionary) algorithm for finding symmetric Nash equilibria in symmetric games. Our implementation is based on the replicator dynamics formalism introduced by Taylor and Jonker [Taylor and Jonker, 1978] and Schuster and Sigmund [Schuster and Sigmund, 1983] and is described in detail in our earlier work [Reeves et al., 2005]. We have found the method particularly useful for finding mixed-strategy equilibria in many-player games with large strategy spaces, but it does not guarantee to find all equilibria.

In 14 out of these 16 environments, $\mathrm{PP}(F^{SC})$ was verified to be an ϵ-Nash equilibrium for $\epsilon < 0.1$. Twelve have $\epsilon < 0.05$, and in six of these (one U and five E) it was an exact equilibrium. The two worst environments were $U(5,3)$ and $U(7,8)$. In the last case, expected payoff for all-$\mathrm{PP}(F^{SC})$ was 2.67, so ϵ represents about 5% of the value. For no other case did it reach 2%. Moreover, the results are quite insensitive to statistical variation. The $\bar{\epsilon}$ values never exceed ϵ by much, and in every environment for which we produced an equilibrium with replicator dynamics, $\mathrm{PP}(F^{SC})$ appears in this symmetric mixed-strategy profile with substantial if not overwhelming probability.

Overall, we regard this as favorable evidence for the $\mathrm{PP}(F^{SC})$ strategy across the range of market-based scheduling environments. 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, $\mathrm{PP}(F^{SC})$ is clearly not a best-response player in this environment.

7 Strategies for Environments with Substitutes

In the previous sections we focused on the exposure problem when there are complementarities in preferences. We found that strategies based on price prediction can be quite effective in mitigating the problem. In this section we extend our analysis of bidding strategies to the case of substitutable goods. The strategic challenge in this environment is bidding when there are significant own price effects: bidding below willingness-to-pay for the marginal unit may lower the price sufficiently on inframarginal units to be a profitable strategy [Ausubel and Cramton, 2002]. We consider simple demand-reduction strategies as well as a sophisticated approach to predicting own price effects inspired by the success of self-confirming price prediction for environments with complementarities. In the environment with substitutes we study, we find that the simple demand-reduction strategies clearly outperform this price predictor.

To analyze bidding strategies in an SAA game with substitutes, we assume that 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 $v_0 = V$ a uniform upper bound on the marginal value of one unit.

In homogeneous-good environments bidders derive the same value from any bundle of k goods regardless of their labels. The definitions of strategies in this section rely on this assumption, though it would not be difficult to generalize their approaches to apply to environments with a more general type of substitutability. The assumption of homogeneous goods is convenient for computational implementation and analysis, however, we believe that it is not essential to our main results.

7.1 Demand-Reduction Strategy

Consider an SAA game with m auctions, each selling one unit of an identical (homogeneous) good. If all agents follow SB, the outcome is that the bidders for the m most highly valued units win them, at a uniform price equal to the value of the most highly valued losing unit (possibly plus the bid increment). This is virtually equivalent to truth-telling in an $m + 1$st sealed-bid uniform-price auction. Like the truth-telling/sealed-bid case, the all-play-SB outcome is efficient (modulo the bid increment), but it is not an equilibrium. In fact, efficient equilibria in the $m+1$st sealed-bid uniform-price auction do not exist [Ausubel and Cramton, 2002]. To motivate a possibly better strategy, consider the intuition for the non-existence of an efficient equilibrium: if a bidder has a positive probability of influencing price in a situation in which the bidder wins a positive quantity, then the bidder has an incentive to shade her bid in a sealed-bid uniform-price auction. Bid-shading leads to inefficient outcomes. This intuition and the failure of SB motivates considering strategies that suppress demand.[16]

We introduce a relatively simple demand-reduction strategy, DR. Let us modify SB by introducing a parameter $\kappa \in [0, V]$ defining the degree of the agent's demand reduction: the agent bids the ask price on the lth cheapest good as long as it is not winning that good, and its marginal surplus is at least $\kappa(l - 1)$. In other words, the agent considers the goods in order of price, adding the lth good to its bundle until the marginal value v_l drops below the ask price plus $\kappa(l-1)$. We denote a specific demand-reduction strategy in this family by $\mathrm{DR}(\kappa)$. The DR strategy family is a simple way of capturing the intuitions of the demand-reduction literature: bidders should shade their bids, and the amount of shading increases with the number of winning goods [Ausubel and Cramton, 2002].

Formally, define $\mathrm{DR}(\kappa)$'s perceived price of the good with the lth lowest myopically perceived price (defined in Section 3):

$$\rho_l(\boldsymbol{B}) \equiv \begin{cases} \beta_l + \kappa(l-1) & \text{if winning the good} \\ \beta_l + 1 + \kappa(l-1) & \text{otherwise,} \end{cases} \tag{5}$$

where $\boldsymbol{\beta}$ is the vector of current bid prices. Agent $\mathrm{DR}(\kappa)$ plays the perceived-price bidding strategy using $\rho(\boldsymbol{B})$. Note that in our definition of $\rho(\boldsymbol{B})$ for this strategy, we assume that goods with different labels are indistinguishable. We use the subscript l instead of i to emphasize that each good is labeled by its myopic price rank order rather than by the auction selling it.

[16] Note that the sunk-awareness modification of SB we introduced in Section 3.2 to address the exposure problem leads to overbidding, as opposed to bid-shading, in this environment. Using the terminology of Definition 1, the perceived-price vector of a sunk-aware strategy is equal to or below the myopic perceived-price vector used by SB, which results in more aggressive bidding. The perceived price of the demand-reduction strategy we introduce in this section is always at least as high as the myopic perceived-price vector.

7.2 Predicting Own Price Effects

The ability of a single agent to affect final prices is strategically central when goods are substitutes. Therefore, the focus of price prediction in the substitutes case is to model this relationship. Specifically, for the homogeneous-good environment, price predictions take the form of a mapping from purchase sizes (i.e., the agent's chosen demand) to final prices. The main role of this prediction is to guide the agent as to when it is beneficial to refrain from bidding on potentially valuable goods.

The assumption that final prices depend on the number of goods the agent is trying to win implies that the agent's prediction of the final price of good i can no longer be represented by a scalar. Let $\pi_{ik}(\boldsymbol{B})$ be the predicted final price of good i given that the agent tries to win k goods and its information state at the current round is $\boldsymbol{B}$. We can think of the agent's *predicted own-effect prices* as an $m \times m$ matrix, in which the rows are auction labels and the columns are the intended purchase sizes. We define an updating rule for π_{ik}, $i, k \in \{1, \ldots, m\}$, similar to the point price-prediction rule described in Section 4.1. The current price prediction for good i when the agent plans to bid on k goods is the maximum of the initial prediction and the myopically perceived price:

$$\pi_{ik}(\boldsymbol{B}) \equiv \begin{cases} \max(\pi_{ik}(\varnothing), \beta_i) & \text{if winning good } i \\ \max(\pi_{ik}(\varnothing), \beta_i + 1) & \text{otherwise.} \end{cases} \tag{6}$$

There is no apparent 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. Therefore, 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.

In the homogeneous-good environment, agents are indifferent between item subsets of equal sizes. Thus, in our strategy, the agent uses price prediction to determine the number k^* of units to buy, but not to identify specific auctions in which to participate in the current round. Formally,

$$k^* = \arg\max_{k} \max_{|Y|=k} \sigma(Y, \boldsymbol{\pi}_{\cdot k}(\boldsymbol{B})),$$

where $\sigma(Y, \boldsymbol{p})$ is the agent's surplus for goods Y defined by Equation (1), and $|Y|$ refers to the number of goods in set Y.

Given k^*, the choice of goods X^* on which to actually bid is based on the current myopically perceived prices, $\rho(\boldsymbol{B})$ (Equation 2). Using myopically perceived prices ensures that the agent never regrets the composition of its bid set (conditional on size) even if its predicted own-effect prices are wrong.

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

The agent breaks ties as in Definition 1. Given X^*, the agent bids $b_i = \beta_i + 1$ (the ask price) for the $i \in X^*$ that it is not already winning. We call this strategy family the *own-effect price predictor* (OEPP) and denote a specific strategy in this family by $\mathrm{OEPP}(\boldsymbol{\pi}^x)$.

Similar to the point price predictor defined for complementary goods, the OEPP family includes SB as a special case when the predicted own-effect prices are a matrix of zeros: $\mathrm{SB} = \mathrm{OEPP}(0)$. As mentioned in Section 7.1, if all players follow SB, the allocation is efficient. Perceived prices based on an own-effect price matrix with positive elements are weakly higher than the myopic perceived prices SB uses. Therefore, an OEPP agent using positive predictions tends to bid on fewer items than is efficient given the others' bids, and never bids on more goods than SB would.

7.3 Self-Confirming Own-Effect Prices

We define the concept of self-confirming own-effect price prediction similarly to self-confirming point price prediction for complementary environments.

Definition 5 (Self-Confirming Own-Effect Prices) *Let Γ be an instance of an SAA game with homogeneous goods. Matrix $\boldsymbol{\pi}$ is a* self-confirming own-effect price matrix for Γ*, if for all $i, k \in \{1, \ldots, m\}$, $\pi_{ik}(\varnothing)$ is equal to the expectation (with respect to the type distribution) of the final price when one agent tries to win k goods and all the other agents follow* $\mathrm{OEPP}(\boldsymbol{\pi})$.

In other words, self-confirming own-effect prices satisfy the condition that if one of the agents bids to win k goods and the other agents "exploit" their own-effect price predictions, that prediction on average *is correct* for all k. We denote the self-confirming own-effect price matrix by $\boldsymbol{\pi}^{SC}$ and the self-confirming own-effect price-prediction strategy by $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$.

To find approximate self-confirming own-effect prices, we follow an iterative procedure similar to that described in Section 5.3. First, we initialize the own-effect predictors with some own-effect price matrix (e.g., all zero) and, sampling from the homogeneous-good type distribution, run many SAA game instances with a profile in which one agent (the *explorer*) ignores its preferences and tries to win a single good, while the others follow OEPP. When average prices obtained by these agents are determined, we replace the first column in the own-effect price matrix with a column vector with all elements equal to the average price, reset the explorer to win two goods and repeat. After the second batch of simulations, we replace all elements in the second column of the own-effect matrix with the average price and increase the explorer's target number of goods by one. We repeat the process, recycling back to a single unit after the exploration target reaches m. When this process reaches a fixed point, we have the matrix of self-confirming own-effect prices, $\boldsymbol{\pi}^{SC}$. We have not investigated whether a fixed point necessarily exists in homogeneous-good environments, but the price predictions converged in this environment within 30 iterations in all of our experiments (see Figure 4).

7.4 Empirical-Game Analysis

We perform analyses similar to, but less extensive than, those reported in Section 6. We analyzed the $m = n = 5$ environment with uniform preferences introduced at

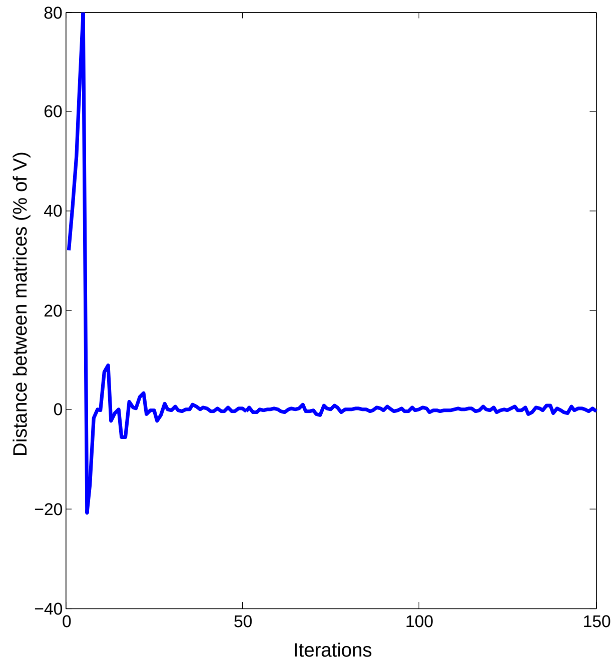

Figure 4: Convergence to a self-confirming own-effect price matrix, starting with an initial prediction that all prices would be zero regardless of the size of the agent's purchase. The prices at each iteration are determined by 10 thousand simulated games. The graph plots the distance between the own-effect prices in consecutive iterations. We define distance between matrices as the maximum over pointwise distances, measured as a percentage of the upper bound on the marginal value, V, of a single unit of the good. The bound V equals 127 in all of our SAA games with substitutes.

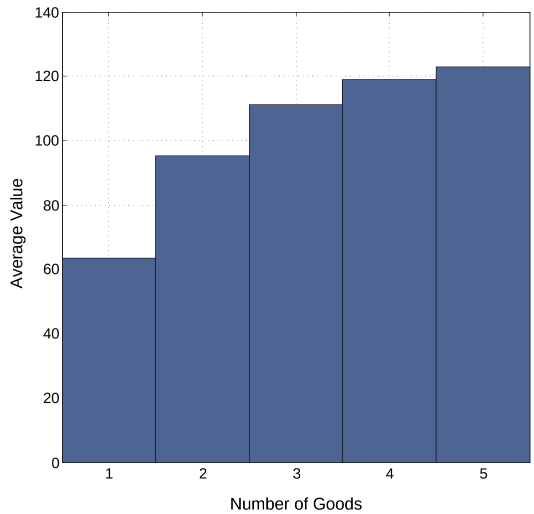

Figure 5: Preference distribution in the homogeneous-good environment.

the beginning of Section 7. We set the upper bound V to 127. In Figure 5 we display the agents' average valuations as a function of the number of goods. As before, our goal is to evaluate the performance of a self-confirming price-prediction strategy, $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$ in this instance. Since the literature predicts that agents suppress demand in equilibrium, we include many instances of our demand-reduction strategy family. We analyzed 51 strategies: SB, 47 $\mathrm{DR}(\kappa)$ with $1 \leq \kappa \leq 120$, one sunk-aware strategy with parameter $k = 0.5$, a self-confirming own-effect price predictor $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$, and the baseline distribution predictor $\mathrm{PP}(F^{\mathrm{SB}})$ (defined in Section 5).

We estimated payoffs for 16542 strategy profiles (out of 3.48 million possible), based on an average of 1.9 million samples per profile. Some profiles are simulated for as few as 40 thousand samples; near-Nash-equilibrium profiles were simulated for up to 205 million game instances per profile. Despite the high-quality information $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$ employs about own effect on final prices, the strategy's use of this information did not provide any advantage over the simpler information-free demand-reduction agents. In the majority of profile settings where it was tested, $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$

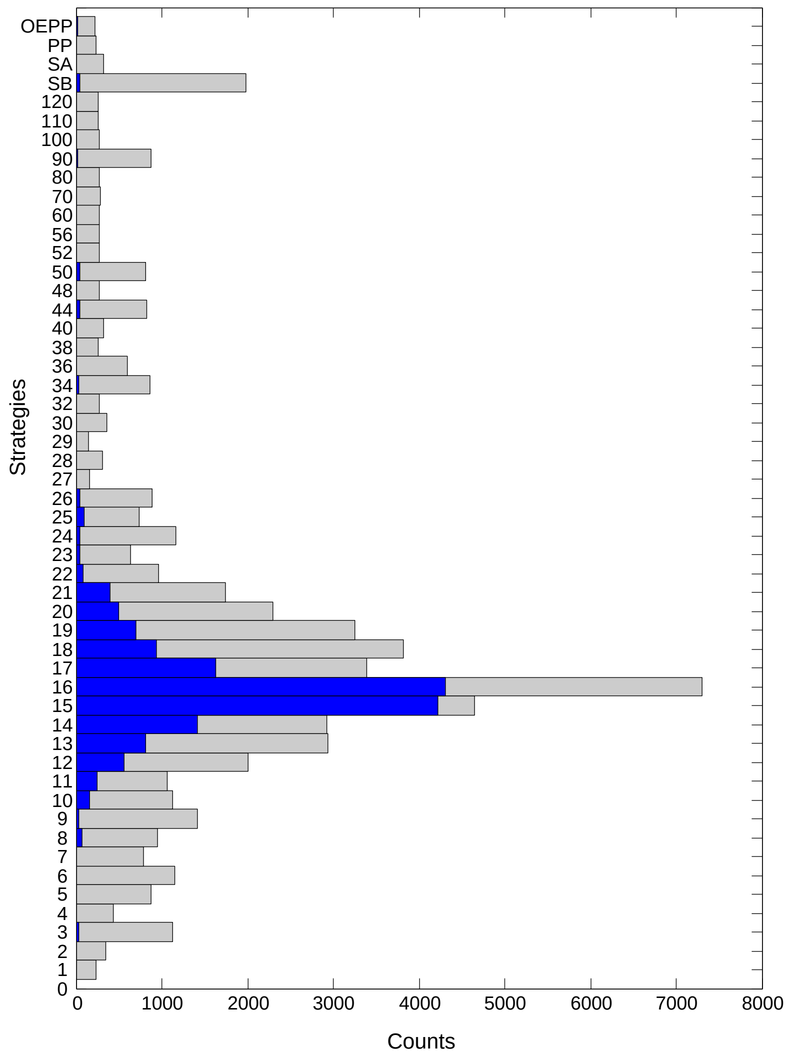

Figure 6: Distribution of best deviations. The light bars reflect the number of estimated profiles in which the corresponding strategy appeared. The dark bars reflect how many times the strategy in fact *was* a best deviation. We index demand-reduction strategies $\mathrm{DR}(\kappa)$ by their corresponding κ-values. OEPP refers to $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$, PP to $\mathrm{PP}(F^{\mathrm{SB}})$, and SA refers to the sunk-aware strategy with $k = 0.5$.

can be refuted with a $\mathrm{DR}(\kappa)$ strategy. Indeed, we found $\mathrm{DR}(\kappa)$ with $10 \leq \kappa \leq 22$ to be the *best deviation* for 96% of 16542 profiles we have analyzed. By the best deviation for a given profile we mean the most profitable of all possible deviations from a strategy in this profile based on neighboring profiles available in our dataset. For 108 profiles our dataset includes estimated payoffs for all deviations from all strategies. The best deviations for these profiles are always $\mathrm{DR}(\kappa)$ with $16 \leq \kappa \leq 19$.

We provide more evidence in Figure 6 by displaying the number of times a strategy was a best deviation (dark bars) relative to the number of estimated profiles in which that strategy appeared (light bars). The latter is proportional to the approximate number of opportunities for that strategy to be a best deviation from some other profile. The dark bars reflect the preponderance of situations in which agents prefer moving toward a $\mathrm{DR}(\kappa)$ strategy with κ near 15. The light bars document our decision, as this evidence was emerging, to focus our finite computational resources on estimating regions of the payoff matrix most important for (near-)equilibrium play.

We found only 14 profiles for which the highest gain can be obtained by deviating to $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$. This is 0.085% of all estimated profiles and 6.36% of all profiles containing at least one $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$ player. We found 40 profiles for which SB is the best deviation. The sunk-aware and $\mathrm{PP}(F^{\mathrm{SB}})$ strategies are never the most attractive deviations in our data.

We found many pure-strategy asymmetric ϵ-Nash equilibria in this environment. Those with the lowest ϵ are profiles of $\mathrm{DR}(\kappa)$ with $14 \leq \kappa \leq 17$. To give a sense of the magnitude of demand (bid) suppression, these κs correspond to 33–40% of the average final unit price if all players follow SB. In Table 3 we present all ϵ-Nash equilibria for which $\epsilon \leq 0.015$[17] and two of our benchmark profiles: all-SB and all-$\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$ (for which the ϵ is rather large). The probability that the profile is an exact Nash equilibrium was estimated empirically as described in Section 6.3. The profiles are listed in the order of increasing ϵ. We have estimated payoffs of all unilateral deviations from the strategies in the near-Nash-equilibrium profiles to all of the other 50 pure strategies. These ϵ-equilibria all consist of $\mathrm{DR}(\kappa)$ with κs in a narrow range; the best deviations are to nearby κs (column 2). If all agents follow $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$, a single agent can improve her payoff by at least 2.86 (5.5% of the average payoff) by deviating to $\mathrm{DR}(24)$.

As expected, equilibrium outcomes are inefficient in this environment. However, the efficiency loss is small: all-16, the symmetric profile with the smallest ϵ, achieves 98.55% efficiency. We present efficiency results for a few symmetric near-Nash-equilibrium profiles and our benchmark profiles in Table 4.

Our results suggest that $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$ is a weak competitor against $\mathrm{DR}(\kappa)$. The weakness of $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$ may lie in its failure to adjust its bidding to its opponents' behavior: having good information does not guarantee strategic advantage. We observe that $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$ bids like an aggressive demand-reduction agent. As a consequence, it earns high profits when playing against other predictors: essentially, in a profile of all-OEPP, players are tacitly colluding to reduce demand and thus prices. Payoffs would be higher if all agents could commit to this behavior. However, when collusion

[17]For reference, the payoffs range from 30 to 69 in our empirical payoff matrix. Thus, the near-equilibrium profiles in Table 3 are quite close to equilibria: the ϵ of 0.015 constitutes at most 0.05% of the payoff.

ϵ-Nash-equilibrium profile	Best deviation	ϵ-gain from one-player deviation	$\bar{\epsilon}$-gain adjusted for sampling error	Probability the profile is exact Nash equilibrium
15 15 16 16 16	15 → 16	0	0.001	0.58
all-16	16 → 15	0.001	0.004	0.25
15 16 16 16 16	16 → 15	0.001	0.006	0.14
15 15 15 16 16	15 → 16	0.004	0.005	0.11
14 16 16 16 16	16 → 15	0.004	0.008	0.02
15 15 15 15 16	15 → 16	0.005	0.009	0
15 16 16 16 17	15 → 16	0.006	0.007	0.11
14 14 15 15 16	14 → 15	0.006	0.008	0.02
14 14 14 16 16	14 → 15	0.007	0.008	0.09
14 15 15 16 16	14 → 16	0.008	0.009	0.02
15 15 17 17 17	17 → 16	0.008	0.012	0
14 14 15 15 15	14 → 15	0.009	0.008	0.07
15 15 15 17 17	17 → 15	0.009	0.010	0.02
14 14 14 15 15	14 → 15	0.010	0.010	0.05
14 15 15 15 16	16 → 15	0.011	0.010	0.02
16 16 16 16 17	17 → 15	0.011	0.010	0.02
all-15	15 → 16	0.012	0.012	0.04
15 15 16 17 17	17 → 16	0.012	0.012	0.01
15 16 16 17 17	17 → 16	0.012	0.013	0
14 14 16 16 16	14 → 15	0.012	0.013	0
15 17 17 17 17	15 → 16	0.012	0.014	0
14 14 14 15 16	16 → 15	0.013	0.015	0
14 15 16 16 16	14 → 15	0.013	0.014	0.01
15 15 16 16 17	17 → 16	0.013	0.013	0
14 14 15 16 16	14 → 15	0.014	0.014	0
all-17	17 → 16	0.014	0.015	0
15 16 17 17 17	15 → 16	0.015	0.015	0
16 17 17 17 17	17 → 16	0.015	0.015	0
all-SB	SB → 14	1.450	1.469	0
all-OEPP	OEPP → 24	2.857	2.905	0

Table 3: ϵ-Nash equilibria for the substitutes environment. The profiles are listed in order of increasing ϵ.

ϵ-Nash-equilibrium profile	Best deviation	ϵ-gain from one-player deviation	Average payoff	Efficiency (%)
all-SB	SB $\to$ 14	1.450	34.266	100
all-14	14 $\to$ 15	0.020	44.665	98.82
all-15	15 $\to$ 16	0.012	45.230	98.69
all-16	16 $\to$ 15	0.001	45.773	98.55
all-17	17 $\to$ 16	0.014	46.307	98.40
all-18	18 $\to$ 17	0.035	46.810	98.26
all-OEPP	OEPP $\to$ 24	2.857	52.063	93.75

Table 4: Efficiency of some symmetric ϵ-Nash equilibria in the substitutes environment. The profiles are listed in order of decreasing efficiency.

is unenforceable, the usual motive to deviate unilaterally is strong.

8 Discussion

Our investigation of bidding strategies for simultaneous auctions leads to qualitatively different conclusions for environments characterized by complementary and substitutable preferences. For the case of complements, we find strong support for a bidding strategy based on probabilistic price prediction, with *self-confirming predictions* derived through an equilibration process. Like other decision-theoretic approaches to bidding [Greenwald and Boyan, 2004], this strategy tackles the exposure problem head-on, by explicitly weighing the risks and benefits of placing bids on alternative bundles, or no bundle at all. The fact that the predictions are self-confirming suggests that this cost-benefit analysis will be accurate when other agents are following the same strategy.

Given the analytic and computational intractability of the SAA game, we evaluated our self-confirming probabilistic price-prediction strategy, $\mathrm{PP}(F^{SC})$, using an empirical game-theoretic methodology. We explored a restricted strategy space including $\mathrm{PP}(F^{SC})$ along with 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 new strategy—very strong evidence in one environment we investigated intensely, and somewhat less categorical evidence for a range of variant environments.

For the case of substitutes, the driving strategic issue is demand reduction rather than exposure risk, and thus it is necessary to predict own price effects as well as exogenous price levels. We defined a bidding strategy, OEPP, based on such predictions, and a concept of self-confirming prices analogous to the approach that proved so successful in complementary environments. In this domain, however, the strategy $\mathrm{OEPP}(\boldsymbol{\pi}^{SC})$ based on explicit self-confirming predictions did not fare well, proving in our empirical experiments significantly inferior to an approach based on simple across-the-board demand reduction.

There are several possible explanations for the relative lack of success of explicit price prediction in substitutes environments. One is that the particular OEPP method we investigated measures own price effects under unrealistic assumptions. Specifically, the strategy predicts the effect of selecting a demand level (number of goods to go for), and sticking with that choice thereafter. In actuality, the agent can and does reconsider its choice at each round conditional on the current auction information. This myopic assumption about the agent's own behavior would tend to overestimate the effect of its immediate decision about demand at the current prices, and thus cause it to reduce demand more aggressively than warranted.

The simple demand-reduction strategy, $\mathrm{DR}(\kappa)$, can pursue an appropriate degree of demand reduction in a particular environment by tuning the free parameter κ. This approach was successful in our experimental environment, but would presumably need to be retuned for a different configuration of goods and preferences. It remains for future work to identify a general approach for deriving robust demand-reduction strategies directly from specification of preference distributions.

Returning to environments with complementarities, our results establish the self-confirming price-prediction strategy as the leading contender for dealing broadly with the exposure problem. 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:

- incorporating price dependencies (but with reasonable computational effort);
- more graceful handling of instances when self-confirming price distributions do not exist;
- more sophisticated prediction updates given price quotes, including possible incorporation of history; and,
- timing of bids: trading off the risk of premature quiescence with the cost of pushing prices up.

Dealing with combinations of complementarity and substitutability, by combining considerations of exposure and demand reduction, is perhaps the most obvious direction for extending the scope of bidding-strategy ideas developed here.

Finally, an indirect contribution of this work is to demonstrate an empirical methodology for game-theoretic analysis when strategy determination is analytically intractable [MacKie-Mason and Wellman, 2005, Wellman, 2006]. We find that even when strategy spaces are enormous, much can be learned by empirically converting an extensive-form game into a normal form in expected payoffs for strategy choices, combined with thoughtful selection of payoff-matrix regions to estimate, and carefully targeted analyses of results.

Acknowledgments

Rahul Suri contributed to the equilibrium analysis in Section 6. This work was supported in part by grant IIS-0414710 from the National Science Foundation.

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.

Lawrence M. Ausubel and Peter Cramton. Demand reduction and inefficiency in multi-unit auctions, 2002. University of Maryland Working Paper 96-07, revised July 2002.

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

Peter Cramton. Simultaneous ascending auctions. In Cramton et al. [2005].

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

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

Jánosz 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.

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

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

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.

Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55:270–296, 2006.

Jeffrey K. MacKie-Mason and Michael P. Wellman. Automated markets and trading agents. In *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.

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

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

Paul Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2003.

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. *Journal of Economic Theory*, 130:220–245, 2006.

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.

P. Schuster and K. Sigmund. Replicator dynamics. *Journal of Theoretical Biology*, 100:533–538, 1983.

P. Taylor and L. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156, 1978.

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

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

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.