Multiplicative Pacing Equilibria in Auction Markets: E-Companion with Proofs and Additional Material

EC.1. Missing Proofs

EC.1.1. Pareto Optimality of Competitive Equilibria

We prove formally that competitive equilibria are Pareto optimal, which is a known result, but we do it specialized to our setting of buyers with budgets for concreteness.

Proposition 7 (The first fundamental theorem of welfare economics). Any competitive equilibrium is Pareto optimal (when considering the utilities of all buyers and the seller).

Proof. For the sake of contradiction, suppose there is a Pareto dominating allocation of goods and money. Let p_j denote the price of good j in the competitive equilibrium, S_i the bundle received by buyer i in the competitive equilibrium (and S'_i in the dominating allocation), and $t_i = p(S_i) \le B_i$ the amount of money buyer i spends in the competitive equilibrium (and $t'_i \le B_i$ in the dominating allocation). Here, we let $p(S) = \sum_{j \in S} p_j$ be the price of any bundle of goods S under the competitive equilibrium prices. Now, for every buyer i, it must be the case that $t'_i \le p(S'_i)$. This is because either $p(S'_i) > B_i$ or $v_i(S_i) - p(S_i) \ge v_i(S'_i) - p(S'_i)$ by the property of competitive equilibria, yet $v_i(S_i) - p(S_i) \le v_i(S'_i) - t'_i$ by Pareto dominance. If $t'_i < p(S'_i)$ for some i, then it follows that the total payment in the dominating allocation is less than the sum of the good prices, contradicting that the seller is at least as well off. On the other hand, if $t'_i = p(S'_i)$ for all i, then no buyer is better off, and also the seller is just as well off, again contradicting Pareto dominance. $\square$

EC.1.2. NP-Hardness of Computing Pacing Equilibria

We first note the following proposition about our Example 1 for modeling binary choices.

PROPOSITION EC.1. *In Example 1, when $\alpha + \delta < 1/3$, no equilibrium satisfies $\min(\alpha_1, \alpha_2) \ge 3\alpha$.*

Proof. The reason is that if such an equilibrium existed, the total price of the first two goods would be at least $6\alpha K_2 = 3(1 - \alpha - \delta)K_1 > 2K_1$ (which follows from the statement). This is the combined budget of the two buyers, resulting in a contradiction. $\square$

With this proposition we are ready to prove our complexity result.

Theorem 3. MAX-REVENUE-PACING, MAX-WELFARE-PACING and MAX-PACED-WELFARE-PACING are NP-complete.

Proof. We reduce an arbitrary 3SAT instance to the following MAX-REVENUE instance. We set T equal to the number of clauses, plus 4 times the number of variables, in the 3SAT instance. For every variable x_j, we create a copy of Example 1, consisting of buyers $1^{x_j}, 2^{x_j}$ and goods $1^{x_j}, 2^{x_j}, 3^{x_j}, 4^{x_j}$, with bids as specified in the example, using $K_1 = 4$, $\alpha = 1/4$, $\delta = 0$, and (hence) $K_2 = 6$. Each of these goods will only be bid on by the buyers corresponding to its own variable (the other buyers have valuation 0 for them). However, the buyers will bid on other goods as well, namely goods corresponding to the clauses. Specifically, we associate buyer 1^{x_j} with the literal $+x_j$, and buyer 2^{x_j} with the literal $-x_j$. A buyer values a clause good at 1 if its literal occurs in that clause, and at 0 otherwise. Finally, we add a single buyer with unlimited budget that values every clause good at 2. Hence, this buyer will necessarily win all the clause goods, at price at most 1 each.

Suppose a satisfying assignment exists. If x_j is set to *true*, set $\alpha_{1^{x_j}} = 1$ and $\alpha_{2^{x_j}} = \alpha$; otherwise, set $\alpha_{1^{x_j}} = \alpha$ and $\alpha_{2^{x_j}} = 1$. This depletes the budgets of the buyers corresponding to variables, resulting in a revenue of 4 times the number of variables. Moreover, for every clause good, the unlimited-budget buyer faces one of the variable buyers with a multiplier of 1, since we had a satisfying assignment. Hence this buyer pays an amount equal to the number of clauses. Hence the MAX-REVENUE-PACING instance has a solution.

Conversely, suppose the MAX-REVENUE-PACING instance has a solution. Then, the unlimited-budget buyer must pay at least an amount equal to the number of clauses. Because she pays at most 1 on each clause good, it follows that she must pay exactly 1 on each clause good. Hence, at least one of the buyers corresponding to positive literals in each clause must have a multiplier 1. But since, by Proposition EC.1, at most one of the two buyers corresponding to a variable can have a multiplier of 1, it follows that these buyers correspond to a satisfying assignment.

Now we switch to the welfare objective. We reduce an arbitrary 3SAT instance to the following MAX-WELFARE instance. We set up buyers corresponding to variables as in the MAX-REVENUE proof. We set $\alpha = \delta = \frac{1}{8}$, $K_1 = 1$, and thus $K_2 = 3$. We let V, C be the sets of variables and clauses in the 3SAT instance, respectively. We set T equal to

$$\delta K_1 + |V|\left(2K_2 + \left(\frac{K_1}{\alpha} + \epsilon\right)\right) = \frac{1}{8} + |V|(14 + \epsilon),$$

For clauses, a buyer values a clause at value $\frac{\delta K_1}{|C|}$ if its literal occurs in that clause, and at 0 otherwise. Finally, we add a single buyer with unlimited budget that values every clause good at $\frac{\delta K_1}{2|C|}$.

Suppose a satisfying assignment exists. Perform the assignment as in the MAX-REVENUE setting. That gives a social welfare of $|V|(2K_2 + (\frac{K_1}{\alpha} + \epsilon))$ from the variable goods. Furthermore, for each clause, at least one satisfied-literal buyer has its pacing multiplier set to 1, thus winning the clause good, yielding utility $\frac{\delta K_1}{|C|}$. Summing over the clauses gives the desired social welfare. Each buyer can at most win all the clauses, and thus their spend is bounded by $(1-\delta)K_1 + \delta K_1$, satisfying their budget constraint.

Conversely, suppose the MAX-WELFARE-PACING instance has a solution. Then each clause good must be allocated to a satisfied-literal buyer. But, in order to beat the unlimited-budget buyer, the satisfied-literal buyer must have a pacing multiplier of at least $\frac{1}{2}$. By Proposition EC.1, this means that the buyer corresponding to the opposite literal must have a multiplier less than or equal to $\frac{3}{8}$. Therefore, the buyers with pacing multipliers of at least $\frac{1}{2}$ correspond to a satisfying assignment.

We can perform almost the same reduction for MAX-PACED-WELFARE-PACING. We construct the same set of buyers and valuations. We set T equal to

$$\delta K_1 + |V|\left(2K_2 + \alpha\left(\frac{K_1}{\alpha} + \epsilon\right)\right) = \frac{1}{8} + |V|\left(7 + \frac{\epsilon}{8}\right),$$

If a satisfying assignment exists, we can set the same pacing assignment as before. The only difference from the previous construction is that the paced welfare from the variable goods is now $|V|(2K_2 + (K_1 + \frac{\epsilon}{8}))$. Combined with the clause good assignment, this gives exactly the desired paced welfare.

The converse case becomes simpler. For any MAX-PACED-WELFARE-PACING solution, it must be the case that each variable has at least one buyer with a pacing multiplier of 1. To obtain the remaining paced welfare of δK_1, these buyers with pacing multiplier 1 must correspond to a satisfying assignment. □

EC.1.3. Correctness of MIP Formulation

Proposition 8. A solution to the MIP (1)-(13) is feasible if and only if it corresponds to the conditions of a pacing equilibrium.

Proof. Assume that all goods j have some buyer i such that $v_{ij} > 0$. Otherwise, we preprocess the problem by removing all goods that no buyers are interested in.

First, let $\alpha_i, x_{ij} \in [0,1], s_{ij} \in \mathbb{R}_+$ be a pacing equilibrium for a pacing game. Let all MIP variables be set according to their definition as it pertains to the pacing equilibrium. Set $x_{ij} = 1$ if $x_{ij} > 0$. If there are multiple buyers with $x_{ij} > 0$ for good j, set $w_{ij} = 1, r_{i'j} = 1$ for two (and only those two) arbitrary buyers $i \neq i'$ among the winners. We now show that all equations are satisfied. Constraint (1) is implied by the third condition of pacing equilibria. Constraint (2) holds since we set $y_i = 1$ exactly when buyer i spends the whole budget. Constraint (3) is implied by our choice of y_i combined with the third condition of pacing equilibria. Constraint (4) is implied by the first condition of pacing equilibria. Constraint (5) is implied by the third condition of pacing equilibria combined with the fact that buyers spend nothing on a good unless they are allocated a non-zero amount. Constraint (6) and (7) are implied by our choice of h_j being the highest bid on good j and the fact that $\bar{v}_j$ upper-bounds v_{ij}. Constraint (8) is implied by our choice for w_{ij}, x_{ij}. Constraint (9) is satisfied because we set p_j equal to the second price, and the constraint is disabled for the highest bid due to $w_{ij} = 1$ and v_{ij} being an upper bound on $\alpha_i v_{ij}$. Constraint (10) is implied by our choice of setting $r_{ij} = 1$ only if buyer i constitutes the second price, and the fact that the constraint is disabled for all other buyers. Constraints (11), (12), and (13) are implied by our choices for w_{ij}, r_{ij}, respectively.

Now assume that we have some satisfying assignment to the MIP. To construct a pacing equilibrium, assign pacing multipliers and spendings according to the values from the MIP, and set $x_{ij} = s_{ij}/p_j$. We now show that each of the three conditions for a pacing equilibrium are satisfied.

Constraint (4) implies $\sum_{i \in N} x_{ij} = \sum_{i \in N} s_{ij}/p_j = 1$. If $x_{ij} > 0$ then $s_{ij} > 0$ and by (5) $x_{ij} = 1$, therefore (6) and (7) imply $\alpha_i v_{ij} = h_j$. For all buyers i' with $x_{i'j} = 0$, we have $\alpha_{i'} v_{i'j} \leq \alpha_i v_{ij}$, otherwise we would violate (6) and thereby contradict our assumption of having a satisfying assignment. This shows that the first condition of a pacing equilibrium is satisfied.

We first show that, in a feasible assignment p_j must be equal to the second price. p_j is both upper and lower-bounded by $\alpha_i v_{ij}$ for the buyer i such that $r_{ij} = 1$. Furthermore, (9) guarantees that p_j is at least as high as the second-highest bid. Finally note that if $\alpha_i v_{ij}$ is the highest bid h_j and $r_{ij} = 1$, then there must exist at least one other buyer such that $\alpha_{i'} v_{i'j} = h_j$ because (13) ensures that $w_{i'j} = 1$ for some i', and (7)-(8) then imply that buyer i' must satisfy $\alpha_{i'} v_{i'j} = h_j$. This shows that p_j is the second price. Now it remains to note that all buyers i with $x_{ij} > 0$ pay p_j, which is exactly the highest bid other than their own for $r_{ij} = 0$. When $r_{ij} = 1$, we established that $w_{i'j} = 1$ for some other buyer, and thus i and i' must be tied for first price, and buyer i is thus still paying the highest bid other than their own. This shows that the second condition of a pacing equilibrium is satisfied.

Constraint (1) ensures that all budgets are satisfied. Constraints (2) and (3) ensure that if budgets are not fully spent then $y_i = 0$, and α_i is then forced to be 1. This shows that the third condition of a pacing equilibrium is satisfied. □

EC.2. Multiple Equilibria with Stochastic Valuations

We provide two examples of pacing games with multiple equilibria where valuations are stochastic.

EC.2.1. Stochastic Pacing Games Equivalence

In this section we show that our deterministic model has a natural stochastic analogue where valuations are not known ahead of time. In this model, we assume that there are n buyers as before. We also assume that there are m categories of goods, and for each category j, buyer i has a weight $w_{ij} \geq 0$ (the weight vector w_i is analogous to the valuation vector v_i of a buyer in a deterministic instance). The buyers participate in ℓ auctions. For each auction, a category j is sampled uniformly at random from the m categories, and a *quality score* $q \in [0, \omega_j]$ is sampled from a distribution with mean 1. The value that buyer i has for the good is then $q \cdot w_{ij}$, where j is the sampled category. Each buyer has some budget B_i. We want to find a vector of pacing multipliers α and a tie-breaking rule $x : [m] \rightarrow [0,1]^n$ specifying how each category is split in case of a tie, such that each buyer's budget constraint holds in expectation. Formally:

$$\ell \cdot \mathbb{E}[x_i(j) p_j] \leq B_i$$

The above stochastic model turns out to be equivalent to a deterministic pacing game where each buyer has a valuation vector w_i, and a budget of $\frac{m}{\ell} B_i$. In particular, pacing equilibria in the deterministic pacing game and in the stochastic model are in a one-to-one correspondence:

PROPOSITION EC.2. *Any vector of pacing multipliers α and corresponding tie-breaking function $x(j)$ is a pacing equilibrium of a stochastic pacing game if and only if it is a pacing equilibrium in the corresponding deterministic game.*

Proof. To show the proposition, it is sufficient to prove that

$$\mathbb{E}\left[x_i(j)p_j\right] = \frac{\ell}{m}\sum_j x_{ij}p_j,$$

since this implies that the pacing complementarity condition holds in the stochastic model if and only if it holds in the deterministic model. To that end, let $k(j)$ be the index of the buyer k with the second-highest bid $\alpha_k q w_{kj}$ for category j. We have

$$\ell \cdot \mathbb{E}_j\left[x_i(j)p_j\right] = \ell \cdot \frac{1}{m}\sum_j x_{ij}\alpha_{k(j)}w_{k(j)j}\mathbb{E}[q] = \frac{\ell}{m}\sum_j x_{ij}\alpha_{k(j)}w_{k(j)j} = \frac{\ell}{m}\sum_j x_{ij}p_j.$$

□

Next, let us see an example of how this reduction can be applied to one of our examples from the paper. We consider a subset of the valuations used in Example 4c shown to the left of Table EC.1. Remember that in this example we have two very different pacing equilibria: $\alpha = (1, 0.01)$ or $\alpha' = (0.01, 1)$.

Table EC.1 Valuations in the infinite goods example. Left: Deterministic valuations. Right: Stochastic valuations.

	g_1	g_2	g_3
v_1	100	1	99
v_2	1	100	99

	$g_1 = 1/3$	$g_2 = 1/3$	$g_3 = 1/3$
v_1	Unif$(0, 200)$	Unif$(0, 2)$	Unif$(0, 198)$
v_2	$\frac{v_1}{100}$	$100 \cdot v_1$	v_1

We can convert this discrete instance into a stochastic model by sampling valuations for v_1 and v_2 in correlated fashion. In particular, applying our reduction from above and sampling the quality score for each category from Unif$(0, 2)$, we get a stochastic setting where we sample columns uniformly as indicated to the right of Table EC.1. Here, v_1 is chosen by uniformly sampling one of three uniform distributions, and then sampling from the chosen distribution. Each uniform distribution corresponds to one of the goods in the deterministic values in Table EC.1, and the expected value of each uniform distribution is equal to the corresponding deterministic value. The valuation v_2 is then a deterministic function of v_1, chosen to preserve the relative relationship between values from Table EC.1. The marginal CDF of v_2 is the same as that of v_1. We sample n goods and set budgets equal to $\frac{n}{3}$. In the stochastic setting we wish to satisfy the budget constraint in expectation. This is achieved by both of the solutions from before, $\alpha = (1, 0.01)$ or $\alpha' = (0.01, 1)$, by linearity of expectation.

EC.2.2. Independently Drawn Valuations

In the second example, we consider the following instance with independently drawn valuations. There are two buyers with budgets $B_1 = 1.2$ and $B_2 = 1$, respectively, and a single good. Each buyer samples their valuation uniformly iid from:

$$v_i \sim \begin{cases} 40 & \frac{1}{2} \\ 4 & \frac{1}{2} \end{cases}.$$

We show two very different pacing solutions that are equilibria. In the first one, both buyers have the same multipliers. Setting $\alpha_1 = \alpha_2 = \alpha$, the total spend is $\alpha\frac{1}{4}(40 + 4 + 2 \cdot 4) = 13\alpha$. It follows that if we set $\alpha = \frac{2.2}{13}$ then exactly $B_1 + B_2 = 2.2$ is spent. A buyer wins when their value is 40 and the other one has value 4, and they may split the good when tied at $(40, 40)$ and $(4, 4)$. In order for budgets to be spent exactly, we solve for the fraction x of the tied goods that buyer 2 should receive as follows:

$$\frac{2.2}{13}(1 + 10x + x) = 1,$$

which yields $x = \frac{54}{121}$.

In the second equilibrium, buyers have different multipliers. We set $\alpha = (1, 0.1)$, in which case buyer 1 wins for every valuation vector except $(4, 40)$; we allocate all of the good to buyer 2 in the case of the valuation vector $(4, 40)$. The spend of buyer 1 is then $\frac{0.1}{4}(40 + 4 + 4) = 1.2$, while the spend of buyer 2 is $\frac{1}{4}4 = 1$.

Buyer 2 achieves expected utility of $\frac{54}{121}(10 + 1) + 10 - 1 \approx 14.9 - 1 \approx 13.9$ in the first equilibrium, but an expected utility of 9 in the second.

The above example was for the case of a single good. However, we can extend this example to multiple goods as follows: say that there are m goods in total, and each buyer draws their valuation v_{ij} iid according to the same distribution as before; i.e.,

$$v_{ij} \sim \begin{cases} 40 & \frac{1}{2} \\ 4 & \frac{1}{2} \end{cases}.$$

Now, we set the budgets equal to $B_1 = 1.2m$ and $B_2 = m$, respectively. By linearity of expectation, the equilibria from the single-good case still cause each buyer to exactly spend their budget in expectation, and thus are equilibria for the m-good case as well.

EC.3. Iterating Best Responses Cycles

Here, we provide step-by-step details of Example 2 which shows that iterating best responses may cycle.

- Initially, buyer 1 wins auctions 1 and 5 and pays 60; buyer 2 wins auctions 2, 3, 4, and 6 and pays 1928. buyer 2 exceeds its budget of 1300 at these multipliers—it exhausts its budget from auction 2 alone, in which it pays 1300, and it also wins three other auctions. Buyer 2's best response is to lower its multiplier so that it wins only auction 2. To do so, buyer 2 sets its multiplier somewhere on the interval $(1300/6503, 5/25) \approx (0.1999, 0.2)$: any lower, and its bid for auction 2 drops below buyer 1's bid of 1300, in which case buyer 2 wins nothing; any higher, and its bid for auction 6 exceeds buyer 3's bid of 5, in which case buyer 2 exceeds its budget.
- After buyer 2 lowers its multiplier, buyer 1 wins more auctions: In addition to what it was winning previously, buyer 1 also wins auction 3 at a price equal to buyer 2's paced bid of at least $300.6(1300/6503) \approx 60.09$. Buyer 1 exhausts its budget of 60 from auction 3 alone. Buyer 1 must set its multiplier low enough to not win auction 3, but such a multiplier is so low that it results in buyer 1 losing all other auctions. Buyer 1's best response is to tie on auction 3, where buyer 2's paced bid is at most $300.6(5/25)$. To do so, buyer 1 sets its multiplier to at most $300.6(5/25)/123 \approx 0.488$.
- After buyer 1 lowers its multiplier, buyer 2 goes from losing to tying on auction 3, causing buyer 2 to pay more than it was previously for that auction, but it also pays much less for auction 2: Instead of paying 1300 for auction 2 as it was previously, it pays around $1300(0.488) = 634.4$. Because buyer 2 is paying so much less for auction 2, it can raise its multiplier to 1, causing it to win auctions 3, 4, and 6 and to pay less than its budget.
- After buyer 2 raises its multiplier to 1, buyer 1 no longer wins auction 3. It can raise its multiplier to 1 and still not exhaust its budget. This brings us back to the first iteration, where all multipliers were set to 1.

EC.4. Experiments on Best-Response Dynamics (One-Shot Setting)

We considered *best-response dynamics* (BR dynamics) to search for a pacing equilibrium in the standard, one-shot setting. We briefly describe these experiments here. BR dynamics can be thought of as a repeated auction market where each buyer has some budget to spend every day and wishes to set its pacing multiplier appropriately. At the end of each day, buyers observe the outcome for the day and best respond to the strategy of the other players. Our goal in these experiments was to see whether warm-starting BR dynamics with the MIP output can improve convergence of BR dynamics and lead it to outcomes with higher welfare than it would otherwise achieve.

We consider two BR algorithms that differ in how the best response is computed. If there is more than one BR pacing multiplier, we break ties towards the highest pacing multiplier (*BR high*), or towards the

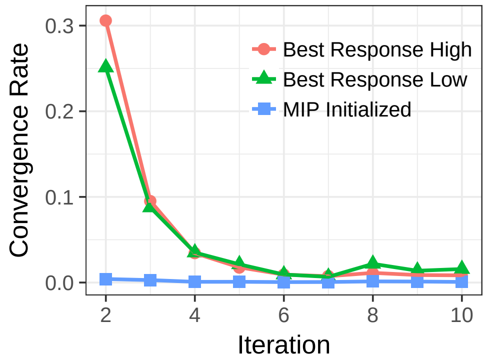

Figure EC.1 BR dynamics convergence rate. For each iteration, we show the absolute difference in a buyer's multipliers from the previous iteration, averaged across buyers and instances. For MIP initialization we average across solutions from all objectives.

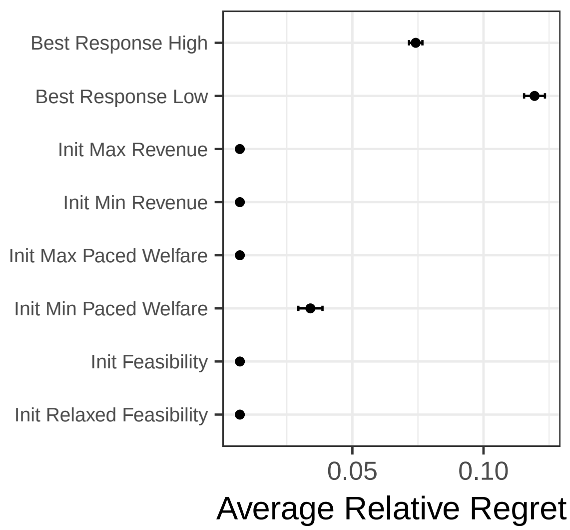

Figure EC.2 Maximum relative regret over buyers in an instance for various BR algorithms averaged across instances.

lowest (*BR low*). Both algorithms always start from the same random initialization of pacing multipliers. In addition, we consider BR high starting from the MIP solutions and refer to it as *Init MIP*. When needed, we replace MIP in the name by a specific MIP objective. For the BR setting, we consider random tiebreaking rather than having fractional allocations be part of the bids. Thus, a pacing equilibrium might not be stable if it includes fractional allocations. We evaluated the BR algorithms on a subset of 50 synthetic instances taken randomly from those in the computational study.

We start by looking at BR dynamics convergence and regret. Figure EC.1 shows that the BR algorithms converge quickly in our computational study. They required less than 10 iterations to reach small oscillations in pacing multipliers. Figure EC.2 shows the maximum relative regret across all buyers, averaged across instances. The relative regret for a buyer is computed as the ratio of the utility-improvement they could get by best responding, divided by the utility of the best response (i.e., the fraction of utility they are missing out on). For the purposes of computing regret, when a buyer exceeds its budget, we do not set utility to negative infinity; instead we penalize utility by the amount over budget multiplied by the spend-to-budget ratio times paced-welfare-to-budget ratio. We see that both BR high and BR low have somewhat high relative regret, missing out on 7.5%-12% utility. Contrary to this, Init MIP solutions perform well and are able to stay near equilibrium for most instances.

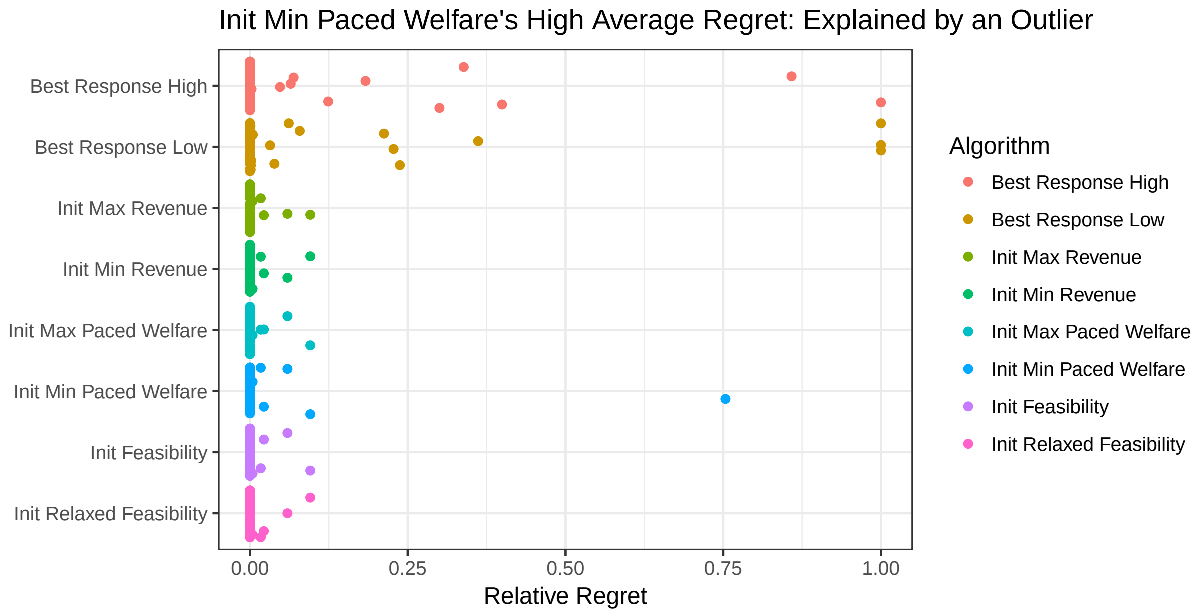

Figure EC.3 Relative regret broken down by each algorithm. Each point represents a BR algorithm running on a particular problem instance.

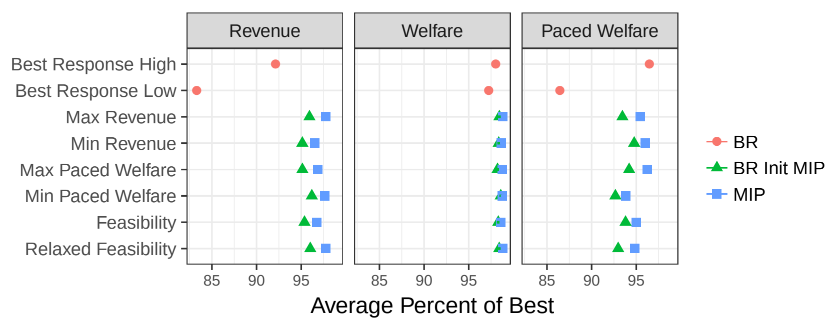

Figure EC.4 Performance of BR and MIP algorithms across 50 instances. The x-axis shows the percentage of the best value for each objective, averaged across buyers and instances.

Figure EC.3 shows the relative regret broken down by each algorithm. This plot shows that the poor performance of initializing with the objective that minimized paced welfare was actually caused by a single outlier.

Finally, we look at the improvement in market outcomes from seeding BR dynamics with the MIP output. Figure EC.4 shows the revenue, welfare, and paced welfare achieved by the different BR dynamics algorithms relative to the MIP. Each point in the plot shows the average performance of a given algorithm relative to the solution maximizing each objective. BR low performs significantly worse than BR high across all three dimensions. For revenue and welfare, they both perform significantly worse than the MIP solutions as well, in spite of the fact that the BR solutions may not even respect budgets. The BR dynamics perform significantly better with MIP initializations than without.

EC.5. Additional Experimental Results and Details

This section describes our experimental setup in more detail than space permitted in the main body of the paper and provides additional results.

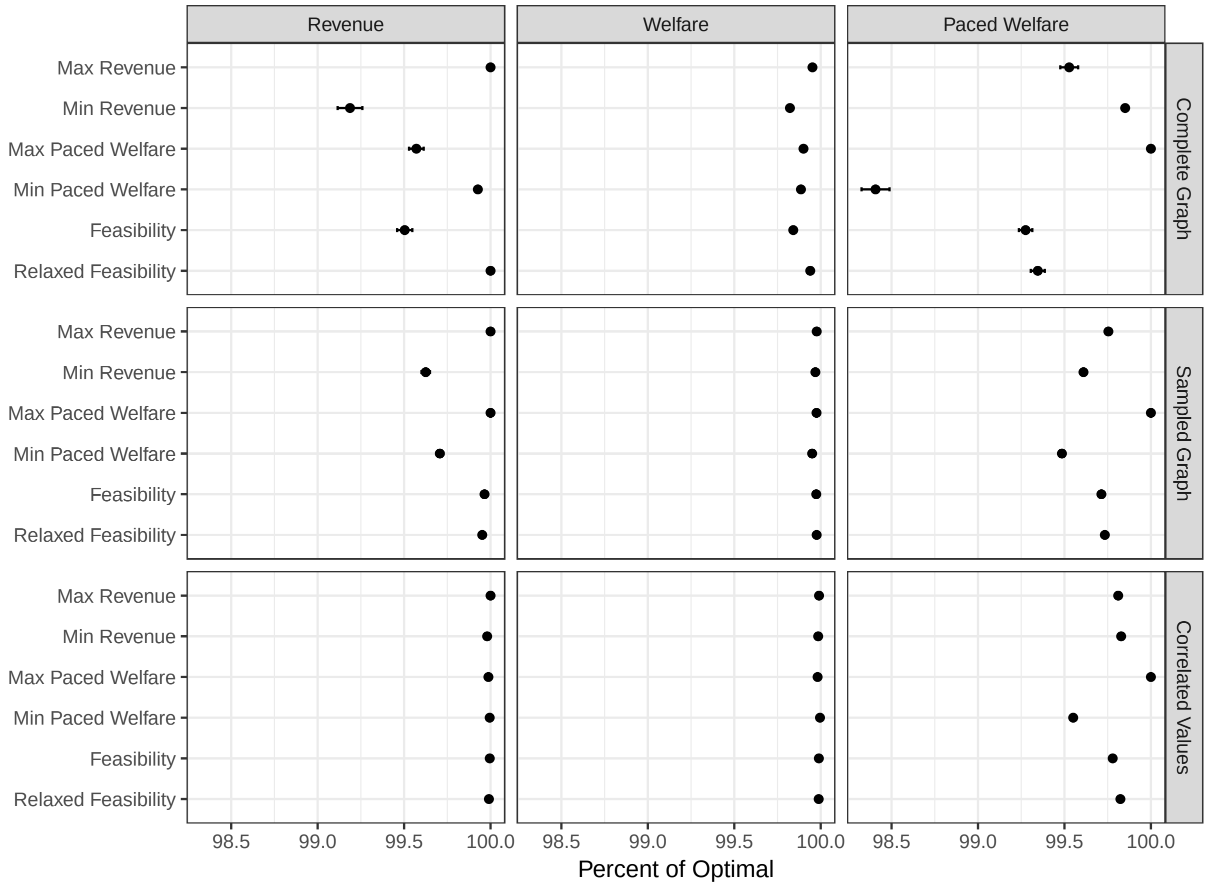

Figure EC.5 Percentage optimality for equilibria maximizing or minimizing the different objectives, measured with respect to each objective, grouped by instance type.

EC.5.1. Equilibrium Gaps

The experiments section reported on maximum gaps for different objectives and instance distributions. Here, we show additional summary statistics. Figure EC.5 shows the relative gap compared to optimal solutions for equilibria maximizing or minimizing the different objectives, measured with respect to each objective, grouped by instance type.

EC.5.2. Other Terms and Notation

We informally defined some terms in the experiments section, which we now define more precisely. For a given scaled-up instance $\tilde{\Gamma}$, let $k_j \in M$ be the *good type* of auction j; this good type associates the auction in the scaled-up instance with a good in the original instance. Let $\tilde{M}_j \subseteq \tilde{M}$ be the set of auctions in the scaled-up instance that have good type j (i.e., $\tilde{M}_j = \{j' \in \tilde{M} : k_{j'} = j\}$ for $j \in M$). For a given run of ADAPTIVEPACING on a scaled-up instance, let x'_{ij} be the *empirical allocation* over good types: the fraction of auctions that buyer i won for good type $j \in M$. That is, let $x'_{ij} = \left(\sum_{j' \in \tilde{M}_j} \tilde{x}_{ij'}\right) / \left(\sum_{i' \in \tilde{N}, j' \in \tilde{M}_j} \tilde{x}_{ij'}\right)$, where $\tilde{x}_{ij'}$ is ADAPTIVEPACING's output allocation ($\forall i \in \tilde{N}, j' \in \tilde{M}$). For a given run of ADAPTIVEPACING on a scaled-up instance, let a buyer's *regret* be the difference between the buyer's maximum possible utility in hindsight (given fixed bids of other buyers) and the buyer's realized utility; let the *max regret* be the maximum such regret across all buyers.

EC.6. Pacing Dynamics

While there are no dynamics in the definition of our game, we consider dynamics to evaluate the quality of the solutions provided by the equilibrium concept. It is instructive to consider the definition of pacing equilibrium in the context of dynamics. Specifically, suppose that the goods are sold continuously over the period $[0,1]$. I.e., at time $t \in [0,1]$ a fraction t of every good will have been sold. Within each infinitesimal slice of time a second price auction is used for each infinitesimal fraction of a good; if there is a tie for a good then it may be split into arbitrary fractions x_{ijt} among the buyers, summing to 1 if there are positive bids. In an ad auction, this would correspond to the limit case where there are large numbers of all types of impressions, and the distribution of such types does not vary over time. Then, we can consider α_i to change dynamically over time (so we get α_{it}). Specifically, if a buyer i is currently spending at a rate that will overspend her remaining budget over the remaining period $[t,1]$, we decrease α_{it}; if it will underspend *and* $\alpha_{it} < 1$, then we increase α_{it}. Call this the *limit dynamics model*.

DEFINITION EC.1. Multipliers $\alpha_i \in [0,1]$ and fractions $x_{ij} \in [0,1]$ constitute a *stable solution* in the limit dynamics model if setting $\alpha_{it} = \alpha_i$ and $x_{ijt} = x_{ij}$ (for all i, j, t) satisfies the feasibility conditions for the x_{ijt} and is consistent with the dynamics (i.e., no α_{it} ever needs to be adjusted up or down).

PROPOSITION EC.3. *Multipliers α_i and fractions x_{ij} constitute a stable solution in the limit dynamics model if and only if they constitute a pacing equilibrium.*

Proof. Suppose they constitute a pacing equilibrium. Then, x_{ijt} is nonzero only if $\alpha_{ijt}v_{ij} = \alpha_{ij}v_{ij}$ is one of the highest bids, and for any j, t, we have $\sum_i x_{ijt} = \sum_i x_{ij} \leq 1$ with equality if there is at least one positive bid. For a buyer with $\sum_j s_{ij} = B_i$ in the pacing equilibrium, we also have $\int_{t=0}^{1} \sum_j s_{ijt} = 1 \cdot \sum_j s_{ij} = B_i$, so the buyer is always exactly on track to spend her budget and the multiplier need not be adjusted. For a buyer with $\sum_j s_{ij} < B_i$ in the pacing equilibrium we must have $\alpha_i = 1$; we have $\int_{t=0}^{1} \sum_j s_{ijt} = 1 \cdot \sum_j s_{ij} < B_i$, so the buyer is always on track to underspend (which is fine because $\alpha_{it} = \alpha_i = 1$). Hence they constitute a stable solution. Conversely, suppose they constitute a stable solution. Then $\sum_i x_{ij} = \sum_i x_{ij0} \leq 1$ with equality if there is at least one positive bid. We also have $p_j = s_{ij}/x_{ij} = s_{ij0}/x_{ij0} = p_{j0}$ which is the second-highest bid $\alpha_{i0}v_{ij} = \alpha_i v_{ij}$. For any buyer i, $\sum_j s_{ij} = \int_{t=0}^{1} \sum_j s_{ijt} \leq B_i$. Finally, if $\alpha_i < 1$ then $\sum_j s_{ij} = \int_{t=0}^{1} \sum_j s_{ijt} = B_i$ (otherwise the multiplier would be adjusted and we would not have $\alpha_{ijt} = \alpha_{ij}$ for all t). $\square$