Optimal selling rules for repeated transactions.

Ilan Kremer and Andrzej Skrzypacz

March 21, 2002

1 Introduction

In many papers considering the sale of many objects in a sequence of auctions the seller is treated as a passive player. We want to extend the analysis of repeated auctions to allow the seller to choose the auctions mechanism.

We present a bunch of examples that consider:

- equilibria in case the bidders are not colluding: i.e. they consider only unilateral deviations
- equilibria with collusive bidders that can coordinate their actions and deviations, but cannot write binding contracts among themselves.

2 The model

There are $N+1$ players: N bidders and a seller. The stage game is the following: in the beginning of each period the seller announces an auction mechanism (set of rules) for that period, i.e. the set of bids B and functions of submitted bids $P_i(\mathbf{b})$ and $T_i(\mathbf{b})$ that determine the probability of player i obtaining the object and the payment of player i. Then the bidders learn their private signals v_i that are distributed independently across players and across time according to a common distribution function $F(v)$ over an interval $[V_L, V_H]$. An important feature of our setup is that in the given period the players know their current (and past) realizations of v_i but do not know their future realizations. The strategies of the buyers are the set of bids and the "not bid": $B \cup \{nb\}$. The payoffs of the bidders are:

$$U(v_i, \mathbf{b}) = P_i(\mathbf{b})v_i - T_i(\mathbf{b}) \tag{1}$$

if player i submits a proper bid and they are 0 if the player chooses not to bid $\{nb\}$.

The payoff of the seller is:

$$U_s(\mathbf{b}) = \sum_{i-1}^{N} T_i(\mathbf{b}) \tag{2}$$

which is independent of $P_i(\mathbf{b})$ as the seller does not value the object per se.

This stage game is played repeatedly over time, infinitely many times. The players choose their strategies to maximize normalized expected discounted sum of the payoffs:

$$(1-\delta)\sum_{t=0}^{\infty}\delta^t U \tag{3}$$

All players use a common discount factor δ. We normalize all payoffs by $(1-\delta)$ to express them in average per-period terms (what simplifies the comparison of equilibrium payoffs for different discount factors).

We define the *static optimum auction* to be a set of rules that maximizes the expected payoffs of the seller in a one-shot game. As Mayerson (????) shown, in that setup the optimal auction can be implemented by a second price auction with an optimal reservation price. We denote by $T^s(b)$ and $P^s(b)$ the corresponding rules of this auction.

We introduce some notation:

W^F - is the expected first order statistic of v. This is the expected value of total surplus that can be generated in an auctions and hence it is an upper bound on equilibrium payoff of any of the players.

W^s_s - is the expected payoff of the seller in a static optimum auction

W^s_b- is the expected payoff of a buyer in a static optimum auction

W^0_s - is the expected payoff of the seller in a SPA with no reserve price.

W^0_b- is the expected payoff of a buyer in a SPA with no reserve price.

We have that $W^F \geq W^s_s + NW^s_b$ where the equality holds only when the optimal reserve price is zero. Similarly, $W^F = W^0_s + NW^0_b$ as the second price auction is efficient and with no reserve price the good is always transferred.

Below we use the notation $W(N)$ to denote the payoffs with N bidders.

3 Non-collusive bidders

We start by creating examples of equilibria in which a patient seller is able to extract the whole surplus in case the bidders do not coordinate their strategies.

3.1 Commitment case

First we illustrate that if the seller can commit at time $t = 0$ to a set of rules, there exists a PBE such that the seller obtains W^s_s in the first auction and (almost) W^F from the remaining periods:

Proposition 1 *Suppose that at $t = 0$ the seller can commit to a set of auction rules. Then for any ε exists a PBE equilibrium such that the seller's expected payoff is*

$$(1-\delta)W^s_s + \delta(W^F - \varepsilon) \tag{4}$$

Proof. Call $T = \delta \left(W_b^0 - \frac{\varepsilon}{N} \right)$. Consider the following set of auction rules: In the first auction the object is allocated according to the static optimum (second price) auction with the payments augmented by T (players are allowed to submit bids $b \geq r$ and $b = 0$). All future auctions are SPA with no reserve price and again with payments augmented by T. If a player does not participate in some auction, he is not allowed to submit any future bids. Given these rules the buyers follow the following strategies: in the first auction bid $b = v$ if $v \geq r$ and bid $b = 0$ otherwise. In all future auctions bid $b = v$.

These strategies are indeed an a best response: The payments T do not depend on b so they only affect the decision whether to bid in the auction or to choose action $\{nb\}$. Conditional on bidding $b = v$ is clearly optimal: the standard reasoning for SPA applies. Also, note that in any stage the payoffs conditional on bidding are at least $\delta \frac{\varepsilon}{N}$ and conditional on not bidding they are 0, so the players are always strictly better off participating in the auctions. ■

The way this selling mechanism works is basically by making the bidders pay an entry fee T for the right to participate in the next auction. This is not a surprising result - we know that if the buyers can be charged a participation fee before they know their values (i.e. the interim participation constraint can be violated) then the seller can extract the whole surplus.

Given that in the first auction the buyers know their values, the most that the seller can extract is:

Lemma 2 *The payoffs of the seller are bounded by*

$$(1 - \delta) W_s^s + \delta W^F \tag{5}$$

Proof. Given the interim participation in the first auction the most that the seller can extract in that auction is W_s^s. The most he can extract in the all future auctions is W^F ■

So the mechanism described in the Proposition above is optimal.

3.2 No commitment

One weakness of the result above is that we have assumed that the seller can credibly exclude forever any player that fails to submit a bid and pay the entry fee. In particular, consider a subgame after all players failed to submit any bids. The seller is then supposed not to hold any auctions. But that strategy is strictly dominated by switching to the static optimal auction forever. So what payoff the seller can extract in an equilibrium if he can only use auction rules that are time-consistent?

The following two examples shed some light on how much revenue the seller may extract:

Proposition 3 *If the optimal auction has no reserve price, then for any $\varepsilon > 0$ for δ large enough so that*

$$(1 - \delta) W_s^s(N - 1) + \delta(W^F(N - 1)) > W_s^s(N)$$

exists a PBE equilibrium such that the seller's expected payoff is at least:

$$(1-\delta)W_s^s(N)+\delta(W^F(N)-\varepsilon) \tag{6}$$

Proof. We construct an equilibrium that obtains this bound. The strategy of the seller is: along the equilibrium path run a SPA with no reserve price with the payments augmented by $T(N)=\delta\left(W_b^s(N)-\frac{\varepsilon}{N}\right)$. If at least one of the players does not submit any bid (in this or any other auction), then one of the non-participants is chosen randomly to be 'excluded' in the next period. An excluded player is not allowed to participate in the current auction. In any period if one player is currently excluded then the seller runs a second price auction among the rest of the bidders with no reserve price and the payments augmented by $T(N-1)=\delta\left(W_b^0(N-1)-\frac{\varepsilon}{N-1}\right)$. If in that auction at least one of the players fails to submit a bid, then one of them is chosen randomly to replace the excluded player (if all submit bids, the excluded player remains excluded in the next period). Finally, if the seller in the history ever failed to follow his equilibrium strategy then he selects the static optimum auction in every period. The equilibrium strategy for all bidders is to bid $b=v$ in all auctions.

First let's verify that following these strategies results in the claimed payoff. In any subgame along the equilibrium path the seller obtains

$$\begin{aligned}W &= W_s^s(N)+NT(N)=W_s^s(N)+N\delta\left(W_b^s(N)-\frac{\varepsilon}{N}\right)\\ &= (1-\delta)W_s^s(N)+\delta(W_s^s(N)+NW_b^s(N)-\varepsilon)\\ &= (1-\delta)W_s^s(N)+\delta(W^F(N)-\varepsilon)\end{aligned}$$

To verify that it is an equilibrium notice that given that all players other than i plan to submit bids in every auction, it is strictly optimal for the player i to do so as well: not bidding gives zero payoff and bidding 0 gives expected payoff of $\delta\frac{\varepsilon}{N}$.

Consider the seller. If he deviates from his strategy, the continuation payoffs are independent from current actions (both for him and for the buyers), so the best deviation is to the static optimum auction. Along the equilibrium path after that deviation his current expected payoff and the expected continuation payoff is lower than from following the equilibrium strategy.[1] Outside the equilibrium path if the seller ever deviated following the strategy is an equilibrium as this is playing the static Nash every period. If the seller always followed the equilibrium strategy but at least one of the sellers deviated the seller if follows the equilibrium strategy obtains:

$$(1-\delta)W_s^s(N-1)+\delta(W^F(N-1)-\varepsilon)$$

[1] Note that unlike in typical repeated games we do not have hear the stadard tradeoff between current gan from deviation and the future punishment. Give that in the stage game the buyers observe the sellers announcement before they submit bids, the punishment is immediate.

and if he deviates he obtains $W_s^s(N)$. By assumption he does not want to deviate. ■

Notice that there always exists $\delta_c < 1$ s.t. for all $\delta > \delta_c$ the condition is satisfied: $W^F(N-1) > W_s^0(N)$ so even if $N = 2$ (and hence $W_s^s(N-1) = 0$) it is satisfied for large δ. In general it is quite mild: it requires that the seller prefers an auction with $N-1$ players plus the expected first order statistic from $N-1$ players in all future auctions over just static auction with N players.

What if the optimal static auction has a non-trivial reserve price?

Proposition 4 *Suppose that the static optimal auction has a non-trivial reserve price. Then if:*

$$(1-\delta)W_s^0(N-1) + \delta W^F(N-1) > W_s^s(N)$$

there exist equilibria with the seller obtaining payoff arbitrarily close to

$$(1-\delta)W_s^s(N) + \delta W^F(N)$$

Proof. The equilibrium is almost the same as before: along the equilibrium path the seller runs the first auction as a static optimum and all future auctions as second price auctions with no reserve price. In all those auctions the payments are augmented by $T(N)$. If any player deviates by not submitting a bid, he becomes excluded. If the seller deviates the continuation equilibrium is always playing static optimum. The bidders bid in every auction $b = v$.

If those strategies are followed, the payoff is stated in the proposition. The incentives of all players are the same as in the pervious proof (the condition assures that the seller has incentives to follow the punishment phase. given that he does not wan to deviate in the punishment phase, he has even less incentives to deviate along the equilibrium path) ■

In case the reserve price is non-trivial the condition in this proposition is satisfied for high δ if:

$$W^F(N-1) > W_s^s(N)$$

It is is satisfied for all N if v is distributed uniformly. From Bulow and Klemperer (1996) we know that $W_s^s(N)$ is smaller than a second order statistic from $N+1$ variables. So for this condition to hold it is sufficient that:

$$V^{(1)}(N-1) \geq V^{(2)}(N+1) \tag{7}$$

where $V^{(k)}(N)$ denotes $k-$ *order* statistic from N draws. Clearly for large enough N that is satisfied, but for small N for some distributions it can be reversed.

In both propositions we have constructed the punishment phase to be such that the buyer once excluded remains excluded forever. We can clearly make the punishment less severe: for example allow him to return to the auction if he bids and pays additional 'entry fee' of $\frac{T(N)}{\delta^k}$ - where k is the number of periods that passed since he deviated, so that the expected present value of his entry fee calculated using the time of deviation as the base period remains unchanged.

If this condition is not satisfied, in order to extract large surplus the seller needs to use less drastic punishments. We show now how a patient seller can again extract the full surplus:

Proposition 5 *For all $F(v)$ all N and all $\varepsilon > 0$ there exists δ_c s.t. for all $\delta \geq \delta_c$ the seller's payoffs are arbitrarily close to:*

$$(1-\delta)W_s^s(N) + \delta W^F(N) \tag{8}$$

Proof. We propose the following equilibrium:

As usual, in the first period the seller runs the optimal static auction with payments augmented by $T(N) = \delta\left(W_b^0(N) - \frac{\varepsilon}{N}\right)$. As long as there is no deviation he announces a second price auction with no reserve price and payments augmented by $T(N)$ in all future periods.

If any bidder fails to submit a bid, the seller starts a punishment phase: he does not run any auctions for k periods. After that he returns to running the second price auctions with no reserve price and augmented payments. If somebody deviates, he restarts the punishment phase. Finally, if the seller ever deviates, he switches to the static optimum forever. The buyer strategy is to bid $b = v$ in all auctions.

To see that for large δ it is an equilibrium first consider the buyers. As above, conditional on bidding they are playing a best response. If they follow the equilibrium strategy they get expected continuation payoff of at least:

$$\delta\left(\frac{\varepsilon}{N}\right) \tag{9}$$

If a bidder deviates his expected payoff is:

$$\delta^k\left((1-\delta)W_b^0(N) + \delta\left(\frac{\varepsilon}{N}\right)\right)$$

A buyer hence have no incentive to deviate if

$$\frac{\delta(1-\delta^k)}{1-\delta}\left(\frac{\varepsilon}{N}\right) - \delta^k W_b^0(N) > 0$$

as $\lim\limits_{\delta\to 1} \delta\frac{1-\delta^k}{1-\delta} = k$ we can clearly find k and δ_c large enough so that for all $\delta \geq \delta_c$ this condition is satisfied.

To check that the seller is playing a best response then along the equilibrium path in every subgame his payoff is at least:

$$(1-\delta)W_s^0(N) + \delta\left(W^F(N) - \varepsilon\right)$$

which for small ε and large δ is clearly higher than $W_s^s(N)$ which is his payoff after deviating. Now consider the punishment phase. If he follows his equilibrium strategy he gets:

$$\delta^k \left((1-\delta)W_s^0(N) + \delta\left(W^F(N) - \varepsilon\right)\right)$$

if he deviates he gets $W_s^s(N)$ forever. Again, for small ε, for any k there exists δ_c large enough so that for all $\delta > \delta_c$ we have

$$\delta^k \left((1-\delta)W_s^0(N) + \delta\left(W^F(N) - \varepsilon\right)\right) > W_s^s(N)$$

so indeed we have an equilibrium. Note that in the process we have constructed two bounds δ_c - one looking at the incentives of the buyers and another looking at the incentives of the seller. Clearly taking the larger of those two is sufficient for the proposed strategies to form an equilibrium ■

In the proof we have used very harsh punishments (against all the players not only the deviators). Clearly that can be weakened to allow for as good equilibria even for smaller δ. One natural possibility is to exclude only the deviator from the k auctions. Then the sufficient condition on δ is that

$$\delta^k \left((1-\delta)W_b^0(N) + \delta\left(\frac{\varepsilon}{N}\right)\right) < \delta\left(\frac{\varepsilon}{N}\right)$$

as before and

$$\delta^k \left((1-\delta)W_s^0(N) + \delta\left(W^F(N) - \varepsilon\right)\right) + (1-\delta^k)(W_s^0(N-1) + \delta W^F(N-1)) > W_s^s(N)$$

3.3 Detail-free auctions

In all the equilibria above the seller runs auctions that are not detail-free: the calculation of the entry fee $T(N)$ requires knowledge of the distribution of values. Given that we usually think that the seller is less informed than the buyers, that is clearly problematic.

If we want a detail-free auction design, the seller can still improve upon the static optimum:

Proposition 6 *Suppose that the static optimum auction has no reserve price. Then for large δ there exist equilibria in which the seller earns arbitrarily close to:*

$$(1-\delta)W_s^s(N-1) + \delta W^F(N-1) \tag{10}$$

Proof. The equilibrium we construct is a follows: the rules of all auctions along the equilibrium path are the same, but the first one is with N players and all the rest are with $N-1$ players. The rules are: all buyers submit a two-element vector of bids. The buyer with the highest first element wins the object and pays the second highest bid. Additionally, one of players that submitted the lowest second element of the bid vector is excluded in the next auction and all others pay his bid. if the seller ever deviates from the equilibrium strategy he plays static optimum forever. The bidders submit bits $b = (v, \delta W_b^0(N-1))$.

In all subgames other than the first one the seller obtains a payoff

$$\begin{aligned} & W_s^s(N-1) + \delta(N-1)W_b^s(N-1) \quad (11) \\ = \quad & (1-\delta)W_s^s(N-1) + \delta W^F(N-1) \quad (12) \end{aligned}$$

and in the first subgame the expected payoff is even higher as the first auction has all N players.

To see that for high δ this is an equilibrium notice that $W^F(N-1) > W_s^s(N-1)$ so the seller has no incentives to deviate. The bidders clearly choose optimally the first element of the bid, as it does not influence their continuation payoffs and in a one-shot game this is a dominant strategy. The second element is also optimal. In fact it is a strictly dominant strategy: the value of being included in the next auction is $\delta W_b^s(N-1)$. ■

The equilibrium described above explores again the idea of charging the players an entry fee. Given that, however, the seller does not know the distribution of values, he has to auction off the rights to be included in the auctions and in order to have the market reveal the price he has to limit the supply (with excess supply the price is 0). That is why he needs to exclude one player.

We see that this simple mechanism for the seller allows him to extract more than the static optimum auction even if he does not know the distribution of values. That example illustrates how creating a group of privileged biders may increase the seller payoffs as in order to be 'included' bidders will be willing to compete away their future payoffs.

4 Collusive bidders

In all the equilibria constructed in the previous section we have treated the bidders as trying to play independently and have only imposed the restriction that conditional on the behavior of other players any one bidder does not want to deviate unilaterally. But what if the bidders coordinate their actions?

In this section we assume that given the strategy of the seller the bidders pick the best strategy for themselves, but are not able to write any binding agreements between themselves.

4.1 Commitment case

First, note that if the bidders can indeed sustain perfect collusion (they always only pay the reserve price and the good goes to the buyer with the highest v) then in a one-shot game the best the seller can do is just to post a price - the optimal reserve price with players having v distributed as the first order statistic from N draws from $F(v)$. Denote expected payoffs corresponding to this mechanism by W_s^c for the seller and W_b^c for all the buyers.

We start by showing that a seller that can commit to a set of auction rules and can extract asymptotically the whole surplus:

Proposition 7 *If the seller can commit to the auction rules for all periods and the bidders coordinate to play their best collusive response, then the seller can extract at least:*

$$\delta W^F \tag{13}$$

Proof. Consider the following set of auction rules: the seller runs for the players in every period the knockout auction described by McAfee McMillan (1982) augmented by payments $T = \delta\frac{W^F}{N}$. The rules of the knockout auction are the following: all players submit bids. The highest bid wins the object, but his bid instead going to the seller is divided equally among the losing bidders. Given that there is no reserve price the expected total payoff for the bidders (gross of T) is W^F - this auction has an efficient equilibrium and the good is always sold with the seller getting 0 revenue. If any player ever fails to submit any bid, the seller stops selling the objects forever.

Given this strategy of the seller the buyers cannot deviate in any way, either individually or as a group to increase their payoffs. Gross of the fees they are maximizing joint payoff and every player is playing individual best response. Any deviation to $\{nb\}$ implies zero payoffs forever, while sticking to the equilibrium strategies guarantees strictly positive payoff. ■

This strategy for the seller is asymptotically optimal: for large δ he is able to extract almost the whole surplus.

For low δ this is a bad strategy, but for low δ the cartel will have a hard time to provide any incentives for the cartel members to stick to a collusive agreement. Hence the results from the previous section are more appropriate.

4.2 Non-commitment case

Suppose that the seller cannot commit to a strategy at time 0. Instead, as in the previous section, every period he can choose a new auction. Then the threat of no more auctions after the players fail to pay the entry fee is no longer credible.

We consider a case where δ is large enough that if the seller treats all the buyers symmetrically after every history of the game, then they are able to sustain perfect collusion. For any current auction if the buyers decide to participate they can run a knockout auction before the true auction as described in McAfee McMillan (1982) with transfer payments with a threat of switching to competitive bidding forever if somebody deviates.

In that setup we first show that the best the seller can do with symmetric auction rules is just to post a price every period.

Proposition 8 *For any ε there exists $\delta_c < 1$ such that for $\delta > \delta_c$ there exist PBE in which the seller obtains the payoff at least*

$$\delta(W^F - \varepsilon) \tag{14}$$

regardless whether the buyers are colluding or not.

Proof. We suggest the following set of auction rules: along the equilibrium path in every auction the seller runs the knockout auction[2] with all players with the payments augmented by $T = \delta\left(\frac{W^F - \varepsilon}{N}\right)$ (so the total payments he receives is $\delta\left(W^F - \varepsilon\right)$. Along the equilibrium path all the bidders submit bids as they would in a one-shot knockout auction.

If there is a buyer's deviation (at least one of the players does not submit any bids) then the seller does not held any auctions for k periods. If the seller ever deviates, he switches to a static optimal auction (which depends on the conduct of the bidders but with payoff $W_s^s < W^F$).

To see that this is an equilibrium note that conditional on all the players bidding, they play a best response, both as individual buyers and as a cartel. Second, the continuation payoff of each player along the equilibrium path is

$$\delta\frac{\varepsilon}{N} \tag{15}$$

so not deviating yields at least:

$$-(1-\delta)T + \delta\frac{\varepsilon}{N} \tag{16}$$

a deviation (not to submit a bid) gives a total payoff of: at most:

$$\delta^k\left((1-\delta)W^F + \delta\frac{\varepsilon}{N}\right)$$

so for no incentives to deviate for any particular player it is sufficient that:

$$\delta\frac{\varepsilon}{N} - (1-\delta)T - \delta^k\left((1-\delta)W^F + \delta\frac{\varepsilon}{N}\right) > 0 \tag{17}$$

that is equivalent to:

$$\frac{\delta(1-\delta^k)}{1-\delta}\frac{\varepsilon}{N} - T - \delta^k W^F > 0 \tag{18}$$

the limit of that expression as $\delta \to 1$ is

$$k\frac{\varepsilon}{N} - T - W^F \tag{19}$$

clearly for any $\varepsilon > 0$ we can find δ_c and k such that for all $\delta > \delta_c$ the incentive compatibility constraint is satisfied. We can also see that when this constraint holds for a single-player deviations, then it holds also for deviations of the whole cartel.

Finally, to see that the seller is playing a best response in all subgames note that following the strategy he expects in any subgame at least

$$\delta^{k+1}\left(W^F - \varepsilon\right) \tag{20}$$

[2] See the previous proof for the rules of a knockout.

while after deviation he gets

$$W_s^s \tag{21}$$

for any ε small enough so that $W^F - \varepsilon > W_s^s$ (which clearly exists) there exists a δ_c' s.t. for all $\delta > \delta_c'$ the first expression is strictly higher, so the seller is also playing a best response. By taking for the critical δ the maximum of δ_c and δ_c' we satisfy all incentive compatibility constraints and indeed we have an equilibrium ■

Remark 9 *1. Asymptotically in this equilibrium the seller is able to extract all surplus. What is important is that along the equilibrium path he does need to know whether the buyers are able to collude or not.*
2. There are two major problems with this equilibrium First, as $\delta \to 1$ the assumption that the buyers do not know their future v becomes unrealistic. Instead we should assume that they know their values for some time interval T and do not know their values after that. We conjecture that in that case the seller can extract at least close to $e^{-rT}W^F$ if he can run the auctions frequently: he can charge at time t entry fees for time $(t+T)$. Of course, if he knows whether the players are colluding or not, he can extract even more, by running an optimal static auction given the buyers conduct plus charging entry fees for period $t+T$.
3. The second major problem is that now it is not detail-free. And unlike with a competitive conduct, the seller cannot induce a collusive cartel to reveal the entry fee by simply running an auction for the right to participate in the next auction: the cartel will collude to set very low fees.
4. Related problem is that in many auctions we would think that as the frequency of auctions increases, the values become correlated over time. That changes the model a lot, however. But suppose that the current value is correlated with values till time T and uncorrelated with later values. Then the mechanism from point 2 should still work.