Flow Auctions

Ilan Kremer, Zvi Wiener and Eyal Winter†

ABSTRACT.

Flow goods (like electricity) are sold through auctions in a dynamic framework. An important design question is the frequency of such auctions. We use a simple dynamic auction model in continuous time to answer this question. We focus on the relationship between the persistency of bidders' valuations and the optimal choice of frequency. If the seller focuses on the equilibrium in which bidders follow a repeated static Nash then the frequency of auctions should typically increase when persistency declines. However, accounting for the fact that bidders can follow different equilibria that are collusive in nature, the comparative statics is reversed, forcing the seller to reduce the frequency when bidders' valuations are less persistent. The argument builds on the fact that high frequency auctions are more conducive to collusion among bidders.

Key Words: Auctions, Dynamics, Contracting, Nash Equilibrium, Collusion

D86, D44, C72

---

† Ilan Kremer; Stanford University, School of Business 655 Knight Way, Stanford, CA 94305 USA kremer_ilan@GSB.Stanford.Edu. Zvi Wiener; The Hebrew University, School of Business, Jerusalem 91905, Israel Zvi.Wiener@huji.ac.il. Eyal Winter; The Hebrew University, The Center for the Study of Rationality, Jerusalem 91904, Israel. Eyal.winter@huji.ac.l

1. Introduction

In many markets a flow of goods or services is sold through repeated auctions, (e.g., electricity auctions, Treasury auctions). A natural question is how the auctioneer should design such a dynamic scheme. Current auction theory has focuses on the case of a single auction in which bidders compete only once and there are no future interactions. In practice the dynamic aspect is often ignored and the design is a sequence of standard static auctions (e.g. uniform price, second price) where the frequency is determined by efficiency considerations.

Ignoring dynamic strategic considerations seems at odds with the theoretical and empirical findings that agents tend to collude in a repeated game environment. Agents can maintain a non-competitive outcome in equilibrium by using future payoffs as a way to reward or punish each other. This raises the question of how should the auctioneer incorporate such effects into the design. In this paper we take a very basic first step at by examining the choice of the auction frequency. A more ambitious goal that we leave for future research would be to examine more sophisticated dynamic designs in which the auctioneer uses future auctions as a way to reward bidders for aggressive bidding.

A central focus of our analysis will be to examine the relationship between the persistence of bidders' valuations and the optimal choice of auction frequency. In doing so, we will consider the following clear tradeoff faced by the seller. On one hand, efficiency considerations require the frequency to be high (absent of costs to set up the auction). This is because by increasing the frequency we increase the expected length of time in which the good is possessed by the bidder who values it the most, thus increasing efficiency. On the other hand, a high frequency facilitates more convenient conditions for bidders to collude against the seller. If bidders collude on a scheme that keeps bids low and allocate the good to the bidder with the highest valuation, then bidders are more tempted to break the collusion when the frequency is low than when it is high. Under a low frequency contracts involve large packages making the current gain more substantial for the deviator and generating more aggressive bidding.

Taking the intuition above one step further, one might think that in a world without collusion, lower persistence will result in higher frequency. The reason would again be based on efficiency considerations. As bidders' valuations change more rapidly the efficiency gain from making the auction more frequent increases.

We examine a continuous-time infinite-horizon model in which *n* bidders compete for a flow of goods. The auctioneer is using a sequence of first-price auctions and decides on the frequency of these auctions. From the bidders perspective, this is an infinitely repeated game. Hence, there is a continuum of equilibria. We examine the implications of different equilibrium selection rules.

We first examine the benchmark case in which agents play the static Nash equilibrium repeatedly; this can be viewed as "non-collusive" behavior. We show that with no transaction costs the auctioneer's profits are strictly increasing in the frequency of the auction and so he would choose an infinite frequency. We find that when there is some cost for running the auctions, lower persistence does not always lead to higher frequency, although for a wide range of parameters it does.

Next, we consider the case when agents use an efficient collusive scheme that maximizes their profits; this is the worst efficient scheme for the auctioneer in which he receives the lowest prices. We show that such a scheme can be maintained by the agents if the frequency of the auctions is higher than a given threshold. As a result, the auctioneer should choose the highest frequency at which there is no collusion. We also show that threshold is decreasing in the persistence of bidders' types. Hence, we conclude that the optimal frequency of the auction is decreasing in the persistence of types which is the main result of this paper.

We should note that our analysis is conducted under the simplifying assumption that types are common knowledge to the bidders but not to the auctioneer. The reason for this assumption is tractability. While one can solve dynamic games (at least numerically) with private information using the technique in Abreu, Pearce, and Stachetti (1990), this is limited to the case in which types are independent over time. For the case in which there is persistence in types, which is the focus of this paper, we are unaware of a simple solution technique. [1]

Still, one would expect to find similar result in setups with asymmetric information. A higher frequency of auctions still makes the auction more susceptible to collusion as the value of the current auction is lower as compared to future auctions. Also, lower persistence in types facilitates collusive schemes as there is less heterogeneity across bidders. In this case the alternative to a collusive equilibrium is a competitive outcome in which bidders compete away their profits. Hence, bidders would find it more attractive not to deviate from the collusive scheme. These two effects are the driving forces behind our result. To summarize, while the specific comparative statics analysis and the characterization of the optimal contract depends on the details of the auction model, the tradeoff between short and long contracts as described through the intuition provided above is robust to these details. Revealing this important tradeoff is the main purpose of this paper.

Dynamic auctions and collusion in auctions have been discussed by various authors in the past. McAfee and McMillan (1992) compares the outcome of collusion among bidders with and without side payments. Athey and Bagwell (2003) study a model of an infinitely repeated procurement auction with stochastic cost shocks, focusing on information revelation and its effects on price dynamics. Skrzypacz and Hopenhayn (2004) study tacit collusion in repeated auctions focusing on the role of communications and monitoring in sustaining it.

Some attention has also been devoted in the literature to flow auctions of the type we consider here. Ausubel and Cramton (1998) discuss securities auctions arguing that the Vickery auction is superior to both discriminatory auctions and uniform price auctions. Cramton and Stoft (2006) also discuss the shortcomings of uniform price auctions in the electricity market.

[1] A notable exception is a recent important paper by Athey and Bagwell (2008), who obtain a numerical solution for a model with persistence in types.

Our point of departure from the literature discussed above is primarily defined by the objectives of our paper. We are raising a design question regarding the optimal choice of frequency which, as far as we know, has not been discussed in the literature[2]. Moreover, our paper highlights the important role that collusion and market persistency play in the design of dynamic auctions. Finally, another, more technical, point of departure from the existing literature is our modeling of dynamic auction using a continuous-time framework. This feature of the model allows us to provide comparative static results that would be technically intractable in a standard discrete-time framework. We hope that our modeling strategy will turn out to be useful for addressing other issues on dynamic auctions. While our model is phrased in terms of multiple buyers and a single seller, our results all apply to the dual case in which multiple sellers compete by bidding on a selling price for a single buyer. One may argue that frequent auction enhances efficiency in a world without collusion may be precisely the intuition that led the California Power Exchange to run frequent auctions on short term contracts to respond to the volatility in production costs of electricity producers prior to the 2000 electricity crisis in California (see CBO paper (2001)). The startling observation that we will make here is that because producers *can* collude (against the electricity supplier) the comparative static is reversed. Namely, auction should be carried out less frequently when producers' production costs are more volatile.

2. Model

We consider an infinite-horizon model in which $n>2$ bidders compete for a flow of goods. The auctioneer is using a first-price auction and decides on the frequency. If the frequency is set to $1/k$ then the auctions is conducted at dates $t = i*k$ for $i = 0,1,2,...$ The good that is sold at time $t = i*k$ is the flow between $i*k$ and $(i+1)*k$.

Valuations and types: There are n distinct values $x_1 \leq x_2,... \leq x_n$ representing bidders' valuations. At time $t=0$ these valuations are allocated among the n bidders with

[2] We are not adopting a general mechanism design approach here as we wish to maintain the format of flow auctions environments such as in the electricity industry.

probability $1/n!$ for each permutation. We let $\bar{x} \equiv \frac{1}{n}\sum_i x_i$ and assume[3] that $x_{n-1} > \bar{x}$. In subsequent periods due to shocks that occur according to a Poisson process bidders' valuations are reshuffled and a new permutation is realized. The distribution of time intervals between reshufflings is exponential with a persistency parameter $1/a$, i.e., its density function is given by $f(a,t) = ae^{-at}$ and the cumulative distribution is $F(a,t) = 1 - e^{-at}$. We assume for simplicity that values are common knowledge among bidders.

Preferences: Bidders are risk-neutral with discount rate *r*. An agent with an initial type *x* has a present value for a flow of length *k* given by the expected value of the flow over this time interval:

$$X(x,a,k) \equiv \int_0^k \left(xe^{-at} + \bar{x}(1-e^{-at}) \right) e^{-rt} dt$$

Auction format: The auction is a first-price auction in which bidders are restricted to use non-negative bids. Given that types are common knowledge and payoffs are discontinuous, one needs to pick a tie-breaking rule to ensure that equilibrium exists. We assume a tie-breaking rule in which the highest type wins. Hence, in the static Nash equilibrium the bidder with the highest value wins the auction and pays the value of the second-highest type, $X(x_{n-1},a,k)$. An alternative approach is to assume a price grid. The equilibrium that we choose can be viewed as the limit of the discrete case when the grid becomes finer.

3. Analysis

3.1. Static Nash: a benchmark case

An optimistic outcome for the auctioneer is that bidders ignore the dynamic aspects of the game. That is, bidders repeatedly use the static Nash strategies of a one-shot game; clearly, this is also equilibrium in the dynamic game. In this subsection we examine the

[3] This condition implies that the number of bidders is at least three, as we assumed earlier.

implications of such behavior. This serves as a benchmark for the next subsection where we consider other equilibria.

If agents play static Nash strategies then the expected revenues from each auction equal the second-highest value:

$$X(x_{n-1}, a, k)$$

The present values of the expected revenues are given by

$$\Pi(k,a) = \sum_{i=0}^{\infty} e^{-irk} X(x_{n-1}, a, k)$$

By substituting the expression for $X(x_{n-1}, a, k)$ we get

$$\Pi(k,a) = \sum_{i=0}^{\infty} e^{-irk} \int_0^k \left(x_{n-1} e^{-at} + \bar{x}(1 - e^{-at}) \right) e^{-rt} dt$$

Lemma 1: For any level of persistence in types the auctioneer's profit is increasing with the auction frequency *1/k*, that is $\frac{\partial}{\partial k} \Pi(k,a) < 0$

The above Lemma implies that without transaction costs the auctioneer should choose *k* to be as small as possible. The intuition behind this is quite simple. By increasing the auction frequency, efficiency is enhanced. The bidder who wins the auction is more likely to be the one who most values the good. Since the auctioneer extracts most of the surplus it is in his best interest to set a high auction frequency.

One may find the above proposition unsatisfying as the optimal frequency is infinite. Let us consider another model in which the seller incurs some fixed amount *c* to run a single auction. In such a case the auctioneer would balance the benefits from increasing the frequency against the costs that are given by

$$\sum_{i=0}^{\infty} c e^{-rki}$$

Here we would find an interior solution; let $k^*(a,c)$ denote the optimal frequency as a function of persistence and transaction costs. One may conjecture that a high frequency of type changes (high a) implies that we choose a high frequency for the auction (low k); that is, $\frac{d}{da}k^*(a,c)<0$. Such a result would follow if the cross derivative is positive:

$$\frac{\partial^2}{\partial k \partial a}\Pi(k,a)>0$$

Interestingly, this is not necessarily true. While for many parameters it is positive one can show numerically that the cross derivative can also be negative. Consider for example the case in which the persistency is extremely low, i.e., when reshuffling takes place on average every second. Moving from one auction a day to twice a day makes only a negligible gain in revenue because in both schemes the expected value of the good is $\bar{x}$ for almost the entire contract term. Hence, even if we assume that bidders follow a static Nash strategy then in the presence of costs for running an auction it is not obvious that we will want to use a higher frequency for the auction when there is a high frequency of types' changes.

3.2. Collusive schemes

As mentioned before, the static Nash equilibrium described above might be unrealistic. Bidders can follow an equilibrium strategy that yields them much higher profits. In this subsection we examine the implication of bidders selecting an equilibrium that maximizes their profits. We focus on the class of efficient collusive schemes in which the off-equilibrium is static Nash. Hence our model can be described as sequential equilibrium of the following two stage meta-game:

1) First, the auctioneer picks the frequency of the auction, k.
2) Second, bidders pick an equilibrium that provides them the highest aggregate profit. The equilibrium is chosen from the set of efficient equilibria where static Nash is used as punishment in case of deviation.

The restriction to this set of equilibria is done for simplicity. The fact that types are persistent makes it impossible to use numerical methods à la Abreu, Pearce and Stacchetti (1990). One can impose some additional conditions on $\{x_i\}$ that ensure that the restriction of efficient equilibrium or static Nash is done without loss of generality.

Alternatively one can verify that some simple equilibria outside this set also satisfy our main conclusion.

In our model there exists a simple characterization of the set of efficient collusive scheme:

Lemma 2 : (i) A necessary condition for the existence of an efficient collusive scheme is given by

$$\frac{1}{n}\sum_{i=1}^{\infty}(1-e^{-iak})e^{-irk} \geq 1 \tag{1}$$

(ii) If (1) holds then there exists a collusive scheme with zero prices.

To better understand the above condition (1), consider a bidder who initially does not have the highest type at time t = 0. In this case $\frac{1}{n}\sum_{i=1}^{\infty}(1-e^{-iak})e^{-irk}$ represents the present value of cash flow that equals one in each interval in which his type in the beginning of that interval is the highest type and zero otherwise. The key to the existence of such a scheme is that the bidder with the second highest type prefers such risky cash flow to getting one dollar today. Based on the above proposition we can solve for the optimal frequency which we denote by $k(a)$:

$$k(a) = \inf\{k \mid \frac{1}{n}\sum_{i=1}^{\infty}(1-e^{-iak})e^{-irk} \geq 1\}$$

This can be rewritten as:

$$k(a) = \inf\{k \mid \frac{1}{e^{kr}-1} - \frac{1}{e^{k(r+a)}-1} \geq n\} \tag{2}$$

The reason that this is the optimal frequency is that a higher frequency leads to zero profits for the seller. Lower frequency results in a static Nash behavior. In this range as Lemma 1 reveals it is better to choose the highest possible frequency (lowest k). There is no closed form solution for $k(a)$. Nevertheless in the appendix we prove the main result of this paper which is that the optimal frequency for the auction $1/k(a)$ is increasing in types' persistency:

Theorem 1 (i) $k(a)$ is the unique solution to $\frac{1}{e^{kr}-1}-\frac{1}{e^{k(r+a)}-1}=n$, (ii) $k(a)$ is increasing.

3.4 Independent Private Values

Our model, thus far, has assumed a fixed set of values that are randomly assigned to bidders at each reshuffling event. If bidders have independent private values drawn from some distribution, the analysis will not admit a close-form solution, and therefore general results cannot be obtained in this alternative framework. However, we did produce the analysis for this model under a uniform distribution of bidders' values to derive the equations that determine the optimal seller's policy and proceeded numerically to derive the relation between type's persistency and the optimal auction frequency. Qualitatively the results fit those we had in the benchmark model. The main challenge this model presents, and the reason why it cannot be solved analytically is that whether or not collusion is sustainable does not only depend on the model's parameters but also on the realization of bidders' types. For each length of contract k there will be a threshold *T(k)* above which the collusion will break down, unless k is small enough and collusion is sustainable for any realization of types. The seller's revenue will, therefore, be affected by *k* through its effect on *T*. If *k* is small enough and collusion is sustainable (with probability 1), the seller's revenue will be zero. As *k* increases the probability of collusion declines but the seller looses in terms of the bidders' willingness to pay. The optimal value of *k* clearly depends on the frequency of reshuffling *a*, and can be denoted by *k(a)*. We have derived the function *k(a)* numerically for markets with 4, 5 and 6 bidders, and established the same qualitative results as in our benchmark model.
The figure below displays *k(a)* (top), and the profit of the seller (bottom) for r=0.05 and $0.1 \le a \le 3$. (Red corresponds to 4 players, blue to 5 and green to 6 players). As can be seen contracts become longer as the frequency of reshuffling increases.

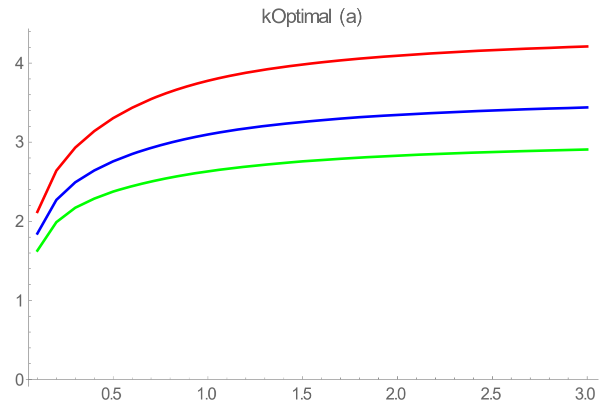

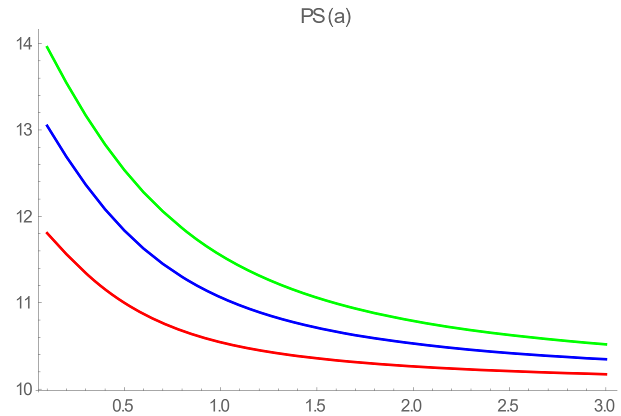

4. Appendix

Proof for Lemma 1: We need to prove that

$\Pi(k,a)=\sum_{i=0}^{\infty} e^{-irk}\int_0^k \left(x_{n-1}e^{-at}+\bar{x}(1-e^{-at})\right)e^{-rt}dt$ decreases with k.

By taking the integral we get

$$\Pi(k,a)=\sum_{i=0}^{\infty} e^{-irk}\left(x_{n-1}\frac{1-e^{-(a+r)k}}{a+r}+\bar{x}\left(\frac{1-e^{-rk}}{r}-\frac{1-e^{-(a+r)k}}{a+r}\right)\right)$$

and by taking the sum this expression can be rewritten as

$\Pi(k,a)=\frac{x_{n-1}-\bar{x}}{a+r}\cdot\frac{e^{rk}-e^{-ak}}{e^{rk}-1}+\frac{\bar{x}}{r}$, in order to study this expression as a function of k we should concentrate on the following (here we use the assumption that $x_{n-1}-\bar{x}>0$):

$\frac{d}{dk}\frac{e^{rk}-e^{-ak}}{e^{rk}-1}=\frac{e^{-ak}\left(-a+e^{rk}(a+r-re^{ak})\right)}{\left(e^{rk}-1\right)^2}$. The sign of this expression is determined by $-a+e^{rk}(a+r-re^{ak})$. This expression can be also written as $-a+e^{rk}(a+r-re^{ak})=(a+r)(e^{rk}-1)+r(1-e^{(r+a)k})$. By replacing the exponents with their Taylor series we obtain

$$(a+r)(e^{rk}-1)+r(1-e^{(r+a)k})=(a+r)\sum_{i=1}^{\infty}\frac{(kr)^i}{i!}-r\sum_{i=1}^{\infty}\frac{(k(a+r))^i}{i!}=$$
$$\sum_{i=1}^{\infty}\frac{rk(a+r)}{i!}\left((kr)^{i-1}-(k(a+r))^{i-1}\right)<0$$
QED

Proof of Lemma 2: consider first the special case in which the price is constant; that is the two highest types bid some price $p\geq 0$. In this case we will argue that (1) is not only necessary but also a sufficient condition. Consider what the incentives of the second highest type from deviating are. According to the collusive scheme his payoff from a single auction in the future, conditional on receiving it (i.e. having the highest value x_n) is:

$$X(x_n,a,k)-p$$

in each interval in which he has the highest type at the beginning of this interval and zero otherwise. His payoff in a static Nash in future intervals is given by:

$$X(x_n,a,k)-X(x_{n-1},a,k)$$

This leads to the difference of:

$$X(x_{n-1},a,k)-p$$

for all auctions starting with the second auction. To get the present value we multiply by the probability of having the highest value and discount to the present.
Hence, we conclude that the benefit from following the collusive scheme as compared to the payoff from deviation is given by:

$$\left(X(x_{n-1},a,k)-p\right)\frac{1}{n}\sum_{i=1}^{\infty}(1-e^{-ika})e^{-rik}$$

We compare this to the current profit from deviation that is given by:

$$X(x_{n-1},a,k)-p$$

Hence, we conclude that (1) is both necessary and sufficient. Since this condition does not depend on *p* (ii) also follows. To complete the proof for part (i) we need to show that if (1) does not hold then we cannot support a scheme with non-negative prices. So we assume by contradiction that $\frac{1}{n}\sum_{i=1}^{\infty}(1-e^{-ika})e^{-rik}=1-\Delta$ for some $\Delta\geq 0$ and that there exist a price scheme in which the price in interval *i* is given by p_i. For this to be an equilibrium it must be that for each *i* we have that:

$$X(x_{n-1},a,k)-p_i\leq\frac{1}{n}\sum_{j=1}^{\infty}(1-e^{-jka})e^{-rjk}\left(X(x_{n-1},a,k)-p_{i+j}\right)\qquad(3)$$

Otherwise at time $t=ik$ there is an incentive to deviate. We consider a weighted average of future prices at time $t=ik$, let

$$\hat{P}_i = \frac{\sum_{j=1}^{\infty}(1-e^{-jka})e^{-rjk} p_{i+j}}{\sum_{j=1}^{\infty}(1-e^{-jka})e^{-rjk}}$$

Using this notation we can rewrite (3) as:

$$X(x_{n-1},a,k) - p_i \le \frac{1}{n}\Big(X(x_{n-1},a,k) - \hat{P}_i\Big)\sum_{j=1}^{\infty}(1-e^{-jka})e^{-rjk}$$

This implies that:

$$X(x_{n-1},a,k) - p_i \le \Big(X(x_{n-1},a,k) - \hat{P}_i\Big)\big(1-\Delta\big) \quad (4)$$

Note that also

$$\hat{P}_i = \alpha p_i + (1-\alpha)\hat{P}_{i+1} \quad (5)$$

where $\alpha = \frac{(1-e^{-ka})e^{-rk}}{\sum_{j=1}^{\infty} e^{-jk(a+r)}}$. Using (4) and (5) we conclude that:

$$X(x_{n-1},a,k) - P_i \le \Big(X(x_{n-1},a,k) - \hat{P}_{i+1}\Big)(\alpha\big(1-\Delta\big) + (1-\alpha))$$

Since $(\alpha\big(1-\Delta\big)+(1-\alpha)) < 1$ we conclude that $\lim_{i\to\infty} X(x_{n-1},a,k) - P_i = \infty$ which implies that $\lim_{i\to\infty} P_i = -\infty$ and we contradict our assumption of non-negative prices. QED

Proof of Theorem 1: Let

$$f(k,a) \equiv \frac{1}{e^{kr}-1} - \frac{1}{e^{k(r+a)}-1}$$

One can easily verify that:

a) f is continuous,
b) $\lim_{k\to 0} f(k,a) = \infty, \lim_{k\to\infty} f(k,a) = 0$, and
c) f is positive.
d) $\frac{\partial}{\partial a} f(k,a) > 0$

We will show that:

e) $\frac{\partial}{\partial k} f(k,a) < 0$

Both parts (i) and (ii) follow from the above properties. To see that e) indeed holds note that:

$$\frac{\partial f(k,a)}{\partial k} = \frac{e^{k(r+a)}(a+r)}{(1-e^{k(r+a)})^2} - \frac{e^{kr} r}{(1-e^{kr})^2}$$

Define $g(x) \equiv \frac{e^{kx} x}{(1-e^{kx})^2}$, we will argue that for every *k*, *g(x)* is a decreasing function of *x*.

We note that

$$g'(x) = \frac{e^{kx}}{(e^{kx}-1)^2} - \frac{2e^{2kx}kx}{(e^{kx}-1)^3} + \frac{e^{kx}kx}{(e^{kx}-1)^2} = \frac{e^{kx}}{(e^{kx}-1)^3}\left(e^{kx} - 1 - 2e^{kx}kx + kx(e^{kx}-1)\right)$$

Denote *y*=*kx*, then the parenthesis can be written as

$$h(y) \equiv e^y - 1 - 2e^y y + y(e^y - 1) = e^y - 1 - e^y y - y$$

To complete the proof we need to show that $h(y) < 0$ for all y>0. To see this note that $h(0) = 0$ and that $h'(y) = -e^y y - 1 < 0$

QED

5. References

Abreu, D., Pearce, D. and Stachetti, E. (1990). "Toward a theory of discounted repeated games with imperfect monitoring," *Econometrica,* 58, 1041-1063.

Athey S. and K. Bagwell (2008) "Collusion with Persistent Cost Shocks", *Econometrica* 76(3) 459-677.

Ausubel, L.M and P. Cramton (1998) "Auctioning Securities", working paper , University of Maryland.

Cramton, P. and Stoft, S. (2007) "Uniform-Price Auctions in Electricity Markets", *Electricity Journal*, 20:1, 26-37.

Congressional Budget Office paper (2001)
http://www.cbo.gov/showdoc.cfm?index=3062&sequence=0

Skrzypacz, A. and H. Hopenhyn (2004) "Tacit collusion in repeated auctions" *Journal of Economic Theory* 114, 1, 153-169.