On the equilibria and efficiency of the GSP mechanism in keyword auctions with externalities

Ioannis Giotis and Anna R. Karlin

University of Washington, Seattle WA 98195, USA**

Abstract. In the increasingly important market of online search advertising, a multitude of parameters affect the performance of advertising campaigns and their ability to attract users' attention enough to produce clicks. Thus far, the majority of the relevant literature assumed an advertisement's probability of receiving a click to be dependent on the advertisement's quality and its position in the sponsored search list, but independent of the other advertisements shown on the same webpage.
We examine a promising new model [1, 16] that incorporates the *externalities* effect based on the probabilistic behavior of a typical user. We focus on the *Generalized Second Price* mechanism used in practice and examine the Nash equilibria of the model. We also investigate the performance of this mechanism under the new model by comparing the efficiency of its equilibria to the optimal efficiency.

1 Introduction

Online search engine advertising is an appealing approach to highly targeted advertising, and is *the* major source of revenue for modern web search engines such as Google, Yahoo! and MSN. The most common setup is as follows: when a user performs a query at a search engine, she is shown a collection of organic search results that contains the links the search engine has deemed relevant to the search, together with a list of *sponsored links*, i.e., paid advertisements. If the user actually clicks on a sponsored link, she will be transferred to the advertiser's web site. For each such click, in which the advertiser receives a potential customer, the advertiser pays the search engine.

Keyword auctions determine which ads get assigned to which keywords (search terms) and how much each advertiser pays. Because of the explosive growth of online advertising and the rising economic importance of ad auctions, a great deal of recent research has focused on developing mathematical models of these systems, with an eye towards understanding their equilibria, dynamics and other properties from the perspective of users, advertisers and search engines [20, 2, 11, 19, 6].

Most keyword auction models assume that each advertisement shown has an inherent click-through rate that depends only on the slot allocated to that

** Email Contact:{giotis,karlin}@cs.washington.edu

advertisement and on the advertisement itself, regardless of the other advertisements that are shown. However, this does not take into account *externalities*: the success of an advertisement depends to a significant extent on which other advertisements are shown alongside it. This is because high-quality competitive ads shown at the same time may detract from each other. Moreover, low-quality ads may deter the viewer from continuing to examine ads shown on the same page.

The importance of the phenomena of externalities has motivated a number of recent papers [12, 1, 16, 9]. Building on work of Craswell et al [8], independently, Kempe and Mahdian [16], and Aggarwal, Feldman, Muthukrishnan and Pal [1] have defined a new Markovian user model. Their model postulates that users scan through the ads in order. For each ad a, users decide probabilistically whether to click, with some ad-specific probability r_a, as well as whether to continue the scanning process, that depends on a different ad-specific probability q_a (as well as the slot the ad is in). This probabilistic continuation models the externality of prematurely terminating the scanning process as a result of either a frustrating irrelevant ad, or a very high-quality web site leading to a purchase. The papers of Kempe et al and Aggarwal et al focus on the problem of computing the efficient allocation given this model, and the efficient (in terms of computational complexity) implementability of an incentive-compatible mechanism (the VCG mechanism).

In this paper, we consider the impact of the new Markovian model on equilibria under the Generalized Second Price (GSP) mechanism. This is important because GSP is the standard mechanism used in practice. Moreover, it is highly unlikely that even if the Markovian model is accurate that the search engines will switch to VCG. Thus, it is of great interest to understand the effects of the new model when GSP is used.

1.1 Results

The focus of our research is on understanding the equilibria of GSP under the new user model. Our main results are the following:

We show that in the new model, in contrast to the most important result about the standard model [20, 11], GSP does not necessarily have an equilibrium in which efficiency is maximized. This raises two key questions: First, does GSP have pure equilibria? And second, how bad can these equilibria be in terms of their efficiency?

The answer to the first question is yes. In Section 3, we give a general construction showing that no matter what the parameters of the system, GSP does have pure Nash equilibria in the new model.

We then turn to the study of the efficiency of GSP equilibria. Our main result here, in Section 4, is that the efficiency of the *worst* Nash equilibrium under GSP can be a factor of k smaller than optimal but no worse, where k is the number of slots in the system. Thus, the so-called *price of anarchy* [18] of GSP with respect to efficiency is k. This latter result depends on the assumption that no advertiser ever bids more than their value. On the other hand, when advertisers can bid more than their value, the efficiency of the worst Nash equilibrium under GSP can be arbitrarily smaller than optimal.

Finally, we show that there are instances where the efficiency of the *best* Nash equilibrium under GSP has efficiency which is a factor of k smaller than optimal. Thus, even the *price of stability* [4] of GSP is k.

2 Model

We consider a model for sponsored search auctions with n participating players (bidders or advertisers) $\{1,\ldots,n\}$, bidding for k advertising slots $\{1,\ldots,k\}$. Each player $i \in \{1,\ldots,n\}$ has three associated values: The first, $0 \leq r_i \leq 1$, represents the *position-independent click-through rate*, which is the probability that the user will click on the ad, given that they look at it. It is a measure of the *relevance* of the ad to the query as well as the general *quality* of the ad, and can also perhaps be thought of as the probability that the user will click on the ad if it is placed in the top slot. The second quantity, $0 \leq q_i \leq 1$ represents the *continuation probability*, the probability that a user will look at the next ad in the list, given that they look at ad i. The third quantity, $V_i \geq 0$ represents the *expected value or profit* of the advertiser given that the user clicks on his ad.

Each advertising slot $s \in \{1,\ldots,k\}$ has an associated fixed constant θ_s representing the ad-independent probability that a user continues scanning advertisements after the s-th slot, given that she scans the s-th slot.

Each player submits a bid and depending on the mechanism used (see discussion below), the search engine produces an allocation of the k slots $\pi()$ such that advertiser $\pi(s)$ is assigned to slot s. An associated list of prices p is also produced such that each time user $\pi(s)$ receives a click he is charged a price of p_s.

We model the behavior of the end-user when presented with the sponsored search results as follows. The user begins scanning the results list with some probability θ_0 which for simplicity we normalize to 1. The first slot is scanned and the user clicks on the ad with probability $r_{\pi(1)}$. Independently of whether the user clicked on the first ad, she proceeds to scan the second slot with probability $\theta_1 \cdot q_{\pi(1)}$, where she clicks on that ad with probability $r_{\pi(2)}$. On the other hand, with probability $1-\theta_1 \cdot q_{\pi(1)}$, the user stops scanning ads after the first and quits the whole process. Given that the user scanned the second slot, she proceeds to scan the third slot with probability $\theta_2 \cdot q_{\pi(2)}$ and so on.

Our main measure for evaluating the performance of the system will the system's *efficiency*. The efficiency of the system for a given ranking of the players π is defined as the sum of the expected utilities of all the players.

$$\begin{aligned} \text{efficiency} &= r_{\pi 1} V_{\pi(1)} + \theta_1 q_{\pi(1)} \left(r_{\pi(2)} V_{\pi(2)} + \theta_2 q_{\pi(2)} \left(\cdots \left(V_{\pi(2)} \right) \right) \right) = \\ &= \sum_{j=1}^{k} \left(\left(\prod_{i=1}^{j-1} \theta_i q_{\pi(i)} \right) \cdot r_{\pi(j)} V_{\pi(j)} \right) . \end{aligned} \tag{1}$$

Discussion

In this model, the probability that the user proceeds to scan the ad in slot $s+1$ given that she scanned the ad in slot s is dependent on both the slot s and the the quality $r_{\pi(s)}$ of the ad in slot s. The dependence on slot has been documented in eye-tracking studies that show that the probability that a user looks at an ad decays with the slot number [13, 15]. This description of the user's eye movement and clicking behavior has been studied under the term "directional market" in several economics papers [5, 3]. As stated, in [3], "the directionality arises due to cognitive burden as it is cognitively 'costlier' for a typical consumer to visit sellers at the bottom of the list before visiting the sellers at the top of the listing".

The dependence of continued scanning on the quality $r_{\pi(s)}$ of the ad in slot s is the combination of two phenomena. First, if the click on slot s results in a conversion, the user is unlikely to continue scanning. Second, if the quality of the ad in slot s is low, the user may be more likely to give up in "disgust". These factors and undoubtedly many others combine to give some ad-dependent probability of continuing to scan. This feature of the model captures the externalities inherent in this setting.

2.1 Mechanisms

The VCG mechanism. One of the mechanisms under examination and our main comparison point is the celebrated Vickrey-Clarkes-Groves (VCG) [21, 7, 14] mechanism. The VCG mechanism is a truthful mechanism which allocates the slots such that efficiency, as defined in (1), is maximized.

Recall that under the VCG mechanism, the expected payment charged to player $\pi'(j)$ at slot j is determined by $OPT_{-\pi'(j)} - (OPT - v_{\pi'(j)})$ where

$$v_{\pi'(j)} = \left(\prod_{i=1}^{j-1} \theta_i q_{\pi(i)}\right) \cdot r_{\pi(j)} V_{\pi'(j)}$$

is the expected utility of this player, OPT is the optimal efficiency with all the players and $OPT_{-\pi'(j)}$ is optimal efficiency without player $\pi'(j)$. Since the most commonly used charging scheme, both in literature and in practice, is on a per click basis, the pay per click price for VCG is defined as

$$p_j = \frac{OPT_{-\pi'(j)} - (OPT - v_{\pi'(j)})}{\left(\prod_{i=1}^{j-1} \theta_i q_{\pi(i)}\right) \cdot r_{\pi(j)}}.$$

The GSP mechanism. Our main focus in this study will be the mechanism most widely used in practice, the *Generalized Second Price* mechanism (GSP):

Definition 1. GSP ***mechanism***

Player	V	r	q	VCG ranking	VCG price(expected)
1	1	1	0.75	1	0.7
2	2	1	0.2	2	0.6
3	0.8	1	0.7	3	0

Fig. 1. Counterexample for the existence of the VCG equilibrium.

Each player i submits a bid b_i representing the maximum amount they are willing to pay for a click. The GSP *mechanism ranks the players in decreasing order of $b_i \cdot r_i$. For the resulting ranking $\pi()$, the price per click of slot j is*

$$p_j = b_{\pi(j+1)} \frac{r_{\pi(i+1)}}{r_{\pi(i)}}.$$

The expected utility $U(\pi(j))$ of player $\pi(j)$ occupying slot j is

$$U(\pi(j)) = \left(\prod_{i=1}^{j-1} \theta_i q_{\pi(i)} \right) \left(r_{\pi(j)} V_{\pi(j)} - b_{\pi(j+1)} r_{\pi(j+1)} \right).$$

We note that in the standard model (where $q_i = 1$ for all i) the GSP ranking maximizes efficiency with respect to the *declared bids*.[1] In our model, this is not the case: if we were to rank by declared efficiency, players with lower $b_i r_i$ might be placed in higher slots than players with higher $b_i r_i$, which would be considered unfair.

3 Nash Equilibria in the GSP mechanism

Of particular interest in the literature on ad auctions [20, 10, 11] has been the equilibrium that yields the same allocation and prices as the VCG mechanism under the standard user model.[2] It is of course very appealing to be able to show that GSP has an equilibrium in which optimal efficiency and several other appealing properties of the VCG equilibrium (such as envy-freeness) hold.

Unfortunately in our model the VCG equilibrium does not always exist. We present a counterexample inspired by a similar counterexample for a different purpose in [1]. Suppose we have 3 bidders and 2 slots with $\theta_1 = 1$. Given the parameters defined in figure 1 it is easy to check that the ranking and VCG prices are as stated in the figure. Notice that the prices in the figure are prices per round or in expectation, therefore the pay per click prices would have to be $p_1 = 0.7$ and $p_2 = \frac{0.6}{\theta_1 q_1} = \frac{0.6}{0.75} = 0.8$ and thus $b_2 = 0.7$ and $b_3 = 0.8$ which cannot result in the desired ranking in the GSP mechanism. Despite the fact the VCG equilibrium might not be achievable, we are able to prove that a pure equilibrium always exists:

[1] Since this mechanism is not truthful, the actual efficiency of the system is not guaranteed to be optimal as in VCG.

[2] We will refer to this equilibrium as the VCG equilibrium.

Theorem 1. GSP *Equilibria Existence*

We assume the players are labeled in decreasing order of $r_i \cdot V_i$. If the players' bids are such that

$$b_s r_s = \begin{cases} V_1 r_1 & \textit{for } s = 1, \\ \sum_{j=s-1}^{k+1} \left(\prod_{i=s}^{j} \theta_{i-1} q_i \right) V_j r_j (1 - \theta_j q_{j+1}) & \textit{for } 1 > s \geq k \\ V_s r_s & \textit{for } k > s, \end{cases} \tag{2}$$

or alternatively by the following recursive definition

$$b_s r_s = \begin{cases} V_s r_s & \textit{for } k > s, \\ (1 - \theta_{s-1} q_s) V_{s-1} r_{s-1} + \theta_{s-1} q_s b_{s+1} r_{s+1} & \textit{for } 1 > s \geq k \\ V_1 r_1 & \textit{for } s = 1, \end{cases} \tag{3}$$

then the resulting allocation and prices of the GSP *mechanism is a Nash equilibrium in the new model.*

The proof of this theorem is presented in the Appendix.

4 The efficiency of GSP equilibria

In light of the fact that the equilibria of GSP may not maximize efficiency, it is interesting to ask how low the relative efficiency (and other properties) of these equilibria can go. We do this using *price of anarchy* and *price of stability* style of analysis.[3] For the price of anarchy analysis we focus on the least efficient GSP equilibrium and compare it against the VCG allocation and the most efficient GSP equilibrium, while, for the price of stability, we compare the most efficient GSP equilibrium against the VCG allocation.

We will also distinguish between two cases. In the first case, the players bid in an unrestricted fashion while in the second case the players can only bid as high as their value. While in reality it is possible for players to bid above their values, it seems unlikely that such bidding behavior can be sustained in practice as the players risk paying a price higher than their value. We therefore expect the restricted case to be more interesting in practice. We will show that the price of anarchy for efficiency can be bounded as per the following theorem.

Theorem 2. *Price of Anarchy*

The price of anarchy of GSP *equilibria both against* VCG *and the best* GSP *equilibrium is k (the number of slots) in the restricted case, and infinite in the unrestricted case.*

Proof. We first look at the efficiency of GSP equilibria in the restricted case. Fix ε and δ arbitrarily small positive constants and consider the following setting. We have $n = k + 1$ players $\{1, 2, \ldots, k, k + 1\}$ bidding for k slots with $\theta_i = 1$ for all $1 \leq i \leq k$. The players' parameters are illustrated in figure 2.

[3] The price of anarchy was originally introduced by Koutsoupias and Papadimitriou in [18] (see also [17] for a survey) as a measure of the performance degradation by selfish autonomous users in the absence of a coordination mechanism.

Player	1	2	3	...	$k-1$	k	$k+1$
V	X	$X-\delta$	$X-1-\delta$	...	$X-1-\delta$	$X-1-\delta$	$X-1-\delta$
q	0	$\frac{1}{1+\delta}$	1	...	1	1	1
r	1	1	1	...	1	1	1

Fig. 2. Example for the price of anarchy regarding efficiency in the restricted case.

It is easy to check that, for large enough X, the most efficient ranking is $[k+1, k, \ldots, 4, 3, 1]$, with total efficiency $kX - (k-1)(1+\delta) \geq kX - \varepsilon X$. Although this ranking is not achievable under GSP, an equilibrium with the ranking $[2, 3, \ldots, k, 1]$ can be achieved if all players bid their values except player 1 who bids $X - 1 - \delta$. The efficiency of this equilibrium is $X - \delta + \frac{1}{1+\delta}\left((k-1)X - (k-2)(1+\delta)\right) \geq \frac{1}{1+\delta}kX - \varepsilon X$ for large enough X.

On the other hand, consider the case under GSP where the players are bidding their values except player 2 who bids $X - 1$. The resulting allocation is $[1, 2, \ldots, k-1, k]$ and we can verify that this results in an equilibrium. Clearly, all players are getting zero utility without being able to improve it. Player 1 is getting utility $X - (X-1) = 1$ while if he were to bid lower to obtain slot j he would still get $1/(1+\delta)(X - (X-1-\delta)) = 1$. The efficiency of this equilibrium is just X. We conclude that for both against VCG and over all GSP equilibria, the price of anarchy regarding efficiency can be bounded by

$$\frac{\frac{1}{1+\delta}kX - \varepsilon X}{X} \geq \frac{1}{1+\delta}k - \varepsilon.$$

We are also able to show that this bound is tight. Indeed, assume we have an arbitrary system of players and slots and consider the least efficient equilibrium of GSP. Focusing on a player x for which $r_x V_x = \max_i r_i V_i$, we will show that the efficiency of this equilibrium is at least $r_x V_x$. Indeed, consider the case where x is not awarded the top slot. In this case some other player y gets the first slot while player x is at slot j with probability of the user getting to that slot ϕ_j. From the equilibrium conditions regarding player x's "desire" to get the first slot by bidding higher than y's bid b_y, we have

$$r_x(V_x - b_y \frac{r_y}{r_x}) \leq \phi_j r_x (V_x - p_j)$$

and using our bidding restriction, we can bound both sides of the inequality.

$$\begin{aligned} r_x V_x - r_y V_y \leq r_x(V_x - b_y \frac{r_y}{r_x}) \leq \phi_j r_x (V_x - p_j) \leq \phi_j r_x V_x \\ r_x V_x - r_y V_y \leq \phi_j r_x V_x \\ r_x V_x \leq \phi_j r_x V_x + r_y V_y \end{aligned}$$

The efficiency of the equilibrium is at least $r_y V_y + \phi_j r_x V_x \geq r_x V_x$. But both the VCG mechanism and most efficient GSP equilibrium cannot have efficiency more than $k \cdot r_x V_x$, hence the price of anarchy is at most k.

For the unrestricted case, consider a setting of 1 slot and two players such that $r_1 V_1 = 0$, $r_2 V_2 = X$. It is easy to see then when $b_1 > X$ and $b_2 = 0$ we have an equilibrium of 0 efficiency. On the other hand, the VCG or optimal GSP equilibrium allocations yield efficiency X. We conclude that for both of these cases the price of anarchy is unbounded.

■

We next turn our attention to the *price of stability* of GSP equilibria relative to the VCG mechanism. Here our goal is to understand how the *best* GSP equilibrium in the worst case compares in performance to the VCG outcome.

Theorem 3. ***Price of Stability***

The price of stability of GSP *equilibria against the* VCG *mechanism is k in the restricted case, and between $k/2$ and k in the unrestricted case.*

The proof of this Theorem is similar in spirit to the proof of Theorem 1 and is omitted from this short version of this paper.

5 Conclusions

We have examined a simple and elegant model for keyword auctions introduced in a series of papers [8, 1, 16] that is able to capture effects that appear in practice but are not considered by the standard model. This model incorporates externalities by modeling the effects advertisements have on the probability that a typical user will scan or click on *other* ads.

Our model makes use of player parameters that are considered a priori determined by the search engine. The use of the click-through probability r_i is generally considered acceptable and these values are probably computed by search engines by sampling the click performance of an ad when the listing is placed randomly in different slots. However, it is not clear if similar techniques can be used to estimate the new parameters q_i. Although determining q_i is not necessary to run the GSP mechanism, if it can be computed efficiently it would certainly open up possibilities for more efficient ranking and pricing schemes.

We have shown that the GSP mechanism always has a pure Nash equilibrium. On the other hand, unlike the standard model, it may not have a Nash equilibrium which maximizes efficiency. We thus attempted to quantify the difference in efficiency between GSP and VCG by examining the price of anarchy and stability. Although the derived bounds appear to make a strong statement in favor of the VCG mechanism, it remains undetermined how these two mechanisms would compare in practice.

An empirical study with real or simulated auction data would potentially reveal more practical results on the performance of GSP and it would be extremely interesting to evaluate its performance against alternative mechanisms that take advantage of the extended information of this model.

References

[1] Aggarwal, G., Feldman, J., Muthukrishman, S., and Pál, M. Sponsored search auction with markovian users. In *Fourth Workshop on Ad Auctions* (2008), ACM Electronic Commerce.

[2] Aggarwal, G., Goel, A., and Motwani, R. Truthful auctions for pricing search keywords. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce* (New York, NY, USA, 2006), ACM, pp. 1–7.

[3] Animesh, A., Viswanathan, S., and Agarwal, R. Competing "creatively" in online markets: Evidence from sponsored search. Working Paper, 2007.

[4] Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., and Roughgarden, T. The price of stability for network design with fair cost allocation. In *FOCS '04: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science* (Washington, DC, USA, 2004), IEEE Computer Society, pp. 295–304.

[5] Arbatskaya, M. Ordered search. Working Paper, 2004.

[6] Cary, M., Das, A., Edelman, B., Giotis, I., Heimerl, K., Karlin, A., Mathieu, C., and Schwarz, M. Greedy bidding strategies for keyword auctions. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce* (New York, NY, USA, 2007), ACM, pp. 262–271.

[7] Clarke, E. H. Multipart pricing of public goods. *Public Choice 11* (1971), 17–33.

[8] Craswell, N., Zoeter, O., Taylor, M., and Ramsey, B. An experimental comparison of click position-bias models. In *1st International Conference on Web Search and Data Mining* (2008).

[9] Das, A., Giotis, I., Karlin, A., and Mathieu, C. On the effects of competing advertisements in keyword auctions. Working Paper, 2008.

[10] Edelman, B., and Ostrovsky, M. Strategic bidder behavior in sponsored search auctions. In *First Workshop on Sponsored Search Auctions* (2005), ACM Electronic Commerce.

[11] Edelman, B., Ostrovsky, M., and Schwarz, M. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. To appear in *American Economic Review*, 2006.

[12] Ghosh, A., and Mahdian, M. Externalities in online advertising. In *Proceedings of the 17th InternationalWorld Wide Web Conference* (2008).

[13] Granka, L. Eye-r: Eye tracking analysis of user behavior in online search. Masters Thesis, Cornell University Library Press, 2004.

[14] Groves, T. Incentives in teams. *Econometrica 41* (1973), 617–631.

[15] Hotchkiss, G. Enquiro eye tracking report ii: Google, msn and yahoo! compared, November 2006.

[16] Kempe, K., and Mahdian, M. A cascade model for externalities in sponsored search. In *Fourth Workshop on Sponsored Search Auctions* (2008), ACM Electronic Commerce.

[17] Koutsoupias, E. Coordination mechanisms for congestion games. *Sigact News 35* (2004), 58–71.

[18] Koutsoupias, E., and Papadimitriou, C. H. Worst-case equilibria. In *16th Annual Symposium on Theoretical Aspects of Computer Science* (1999), pp. 404–413.

[19] Lahaie, S. An analysis of alternative slot auction designs for sponsored search. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce* (New York, NY, USA, 2006), ACM Press, pp. 218–227.

[20] Varian, H. Position auctions. To appear in *International Journal of Industrial Organization*, 2006.

[21] Vickrey, W. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance* (1961), 8–37.

Appendix

A Proof of Theorem 1

Proof. It is easy to check that the two definitions are equivalent. Also, it can be easily seen from the recursive definition that the resulting bids are ordered in the right way as each $b_s r_s$ is a linear combination of $b_{s+1} r_{s+1}$ and the $s-1$'th player's expected valuation. By the initial conditions of the recursive definition it follows that the bids are correctly ordered and the following equilibrium conditions are satisfied

– $\left(\prod_{i=1}^{s-1} \theta_i q_i\right)(V_s r_s - b_{s+1} r_{s+1}) \geq 0$ for all $s \leq k$,
– $\left(\prod_{i=1}^{j-1} \theta_i q_i\right)(r_s V_s - b_j r_j) \leq 0$ for all $j \geq k > s$.

or in other words, all the winning players have greater than or zero utility and the losing players cannot get positive utility by bidding higher.

It remains to show the remaining equilibrium conditions, or that the winning players do not have an incentive to alter their bid so as to get a different slot. Assume an arbitrary winning player s. We need to show that

– For all slots $j < s$, $\left(\prod_{i=j}^{s-1} \theta_i q_i\right)(V_s r_s - b_{s+1} r_{s+1}) \geq V_s r_s - b_j r_j$.
– For all slots $k \geq j > s$, $V_s r_s - b_{s+1} r_{s+1} \geq \left(\prod_{i=s}^{j-1} \theta_i q_{i+1}\right)(V_s r_s - b_{j+1} r_{j+1})$.

To prove the first case, we proceed as follows.

$$
\begin{aligned}
V_s r_s - b_j r_j &= V_s r_s - ((1-\theta_{j-1} q_j) V_{j-1} r_{j-1} + \theta_{j-1} q_j b_{j+1} r_{j+1}) \\
&\quad \text{and since } V_s r_s \leq V_{j-1} r_{j-1} \\
&\leq V_s r_s - ((1-\theta_{j-1} q_j) V_s r_s + \theta_{j-1} q_j b_{j+1} r_{j+1}) \\
&= \theta_{j-1} q_j (V_s r_s - b_{j+1} r_{j+1}) \\
&\leq \ldots \text{ (similarly substituting using the recursive definition)} \\
&\leq \theta_{j-1} q_j \theta_j \cdots q_{s-1} \theta_{s-1} q_s (V_s r_s - b_{s+1} r_{s+1}) \\
&\leq \left(\prod_{i=j}^{s-1} \theta_i q_i\right)(V_s r_s - b_{s+1} r_{s+1}), \text{ since } \theta_{j-1} q_s \leq 1.
\end{aligned}
$$

To prove the second case, for $j > s$, we proceed similarly.

$$\begin{aligned}
V_s r_s - b_{s+1} r_{s+1} &= V_s r_s - ((1-\theta_s q_{s+1}) V_s r_s + \theta_s q_{s+1} b_{s+2} r_{s+2}) \\
&= \theta_s q_{s+1} (V_s r_s - b_{s+2} r_{s+2}) \\
&= \theta_s q_{s+1} \left(V_s r_s - (1-\theta_{s+1} q_{s+2}) V_{s+1} r_{s+1} - \theta_{s+1} q_{s+2} b_{s+3} r_{s+3}\right) \\
&\geq \theta_s q_{s+1} \left(V_s r_s - ((1-\theta_{s+1} q_{s+2}) V_s r_s + \theta_{s+1} q_{s+2} b_{s+3} r_{s+3})\right) \\
&= \theta_s q_{s+1} \theta_{s+1} q_{s+2} (V_s r_s - b_{s+3} r_{s+3}) \\
&\geq \cdots \\
&\geq \theta_s q_{s+1} \cdots \theta_{j-1} q_j (V_s r_s - b_{j+1} r_{j+1}) \\
&\geq \left(\prod_{i=s}^{j-1} \theta_i q_{i+1} \right) (V_s r_s - b_{j+1} r_{j+1}) \, .
\end{aligned}$$

■