

Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions

By BENJAMIN EDELMAN AND MICHAEL SCHWARZ*

In analyses of *sponsored search* auctions for online advertising, it is customary to model the dynamic game of incomplete information by considering a static game of complete information. This approach is used in Edelman, Ostrovsky and Schwarz (2007) (EOS), Varian (2007), and the subsequent literature.

Modeling complex interactions in uncertain environments as games of complete information has a long history. For example, the Bertrand model of oligopolistic competition posits that companies know their competitors' costs based on their experience from prior interactions.

The use of a static game of complete information often offers important benefits. For one, it is tractable – avoiding complex multi-period information sets in a dynamic game. Furthermore, a suitably-chosen static game can capture important characteristics of the underlying dynamic game. When a game is repeated over an extended period, there is good reason to think participants will learn many characteristics of their counterparts – supporting the use of a complete information model.

Yet the use of a static game of complete information is also unsatisfying. Analysis of the static game offers no clear way to identify which (if any) equilibria of the static game of complete information are a relevant approximation of the equilibria of the dynamic game of incomplete information.

In this paper, we consider a dynamic game of incomplete information used to sell sponsored search advertisements. We also consider a corresponding static game of complete information.

* Edelman: Harvard Business School. Schwarz: Yahoo Labs. We thank Susan Athey, Brendan Daley, Robin Lee, Preston McAfee, Michael Ostrovsky, Al Roth, and Adam Szeidl. A draft of this paper was circulated under the title “Optimal Auction Design in a Multi-Unit Environment: The Case of Sponsored Search Auctions.”

We analyze the underlying dynamic game of incomplete information, and we establish an upper bound on the revenue of any equilibrium of any dynamic game in this environment. We then exclude equilibria of the corresponding static game with revenue that exceeds this upper bound. See Section D.

We use this equilibrium selection criteria to assess optimal design of sponsored search platforms that sell search engine advertising. We characterize optimal reserve prices in sponsored search auctions. We show that a search engine's optimal reserve price is independent of the number of bidders and independent of the rate at which click-through rate declines over positions. See Section E.

Our analysis of reserve prices also lets us assess their welfare effects. We separate the effects of reserve price increases into *direct effects* (causing the lowest value bidder to face a higher payment) and *indirect effects* (inducing other bidders to increase their bids, thereby increasing others' payments). We show that most of incremental revenue from setting reserve price optimally comes not from the low bidder's direct effect, and not from indirect effects on other low bidders, but rather from the indirect effects on high bidders. This result may appear counter-intuitive because top bidders' large valuations place them, in an important sense, “furthest from” the reserve price. See Section II.

I. Environment and Mechanisms

A. Environment

Our model for the sponsored search environment follows EOS. We consider a market for advertisements triggered by searches for a single keyword. Each period, a slate of ads is shown to users, and the world lasts for T periods. Each position has a commonly-known click-through

rate (CTR), and higher positions have higher CTR. In expectation, an ad in position i receives α_i clicks per period, and the number of clicks depends on position only. In a given period, each ad can appear in at most one position.

There are N advertisers. Advertiser k values a click at s_k , which is k 's private information, does not change over time, and does not depend on the position where the ad appears. Bidder values are i.i.d. drawn from a commonly-known distribution with support $[0, \bar{s}]$, pdf $f(\cdot)$, and cdf $F(\cdot)$ satisfying the regularity condition from Myerson (1981). In particular, we assume virtual valuation $\psi(s) = s - \frac{1-F(s)}{f(s)}$ is decreasing in s .

The per-period payoff of advertiser k in position i is $\alpha_i(s_k - p_k)$, where p_k denotes the payment per-click made by advertiser k in that period. The total payoff of advertiser k is the sum of k 's period payoffs, *i.e.* $\sum_{t=1}^T \alpha_{k(t)}(s_k - p_k(t))$ where $k(t)$ denotes the position of advertiser k in period t and $p_k(t)$ denotes the payment made by advertiser k in period t .

B. The Dynamic Game of Incomplete Information

We begin with a stylized model that captures the dynamic aspects of the mechanism used in practice. Search engines sell advertisements using real-time generalized second price (GSP) auctions. As each search occurs, an auction is conducted to determine which ads should be displayed to the corresponding user. Advertisers are ranked based on bids: all else equal, the higher the bid, the higher the position assigned to the corresponding advertiser, and hence the more clicks the advertiser receives. If there are more positions than advertisers, a position is allocated to all advertisers whose bid exceeds the reserve price. The per-click payment of the advertiser in the bottom position equals the reserve price. For each other advertiser, the per-click payment equals the bid of an advertiser who is ranked one position lower. (For example, the advertiser in the third position pays an amount equal to the fourth-largest bid.) The mechanisms used by Google, Yahoo!, and Microsoft adCenter build on this GSP approach – though they add variations such as adjusting prices based on ad quality. See e.g. Abrams and Schwarz (2008).

Advertisers can change their bids at any time. Each period, advertisers observe positions of their competitors. By seeing other advertisers' positions, each advertiser updates its beliefs about others' valuations, influencing bids in future periods. When choosing a bid, an advertiser may consider how its action will influence future play of other advertisers. Historic bid information is relevant in equilibrium because an advertiser's best response depends on its expectation of bidding behavior of other advertisers. This information structure creates a complex dynamic game of incomplete information.

Our key observation is that we can establish an upper bound on revenue in any equilibrium of any dynamic game in this environment without characterizing the equilibrium of a dynamic game. Importantly, the *environment* is fairly simple – letting us characterize an optimal mechanism for this environment, thereby bounding revenue in any dynamic mechanism in this environment without solving for an equilibrium of a dynamic game. An optimal mechanism is presented in Section D.

C. Approximation with a One-Shot Game of Complete Information

The dynamic mechanism described above is sufficiently complex that it is difficult or perhaps impossible to solve analytically for an equilibrium strategy. The literature on sponsored search auctions sidesteps the complexity of the dynamic game of incomplete information by modeling it as a one shot game of complete information. The complete information game corresponding to the incomplete information game retains payoffs and the actions exactly as in the previous section. However, there are two important differences: the complete information game lasts only one period, and the per-click values of all bidders are common knowledge. The justification for considering a complete information model is that in a dynamic game of incomplete information, players learn each others' values, making a complete information game a plausible approximation of long-run outcomes. This justification follows longstanding practice in the industrial organization literature, modeling price competition among firms as a one-shot game with common knowledge of each firm's produc-

tion costs.

The one-shot game of complete information has multiple equilibria. Some of these equilibria can be ruled out by the envy-free restriction introduced in EOS and Varian (2007). The envy-free condition requires that no bidder wishes to exchange positions with a bidder ranked above him. Note that the payment of the advertiser in position $i - 1$ equals the bid of advertiser i , so envy freeness means that for each advertiser, $\alpha_i(s_k - p_k) \geq \alpha_{i-1}(s_k - b_k)$ where p_k is advertiser k 's payment per click in his current position, and b_k is k 's current bid. EOS and Varian (2007) show that the equilibrium with VCG-equivalent payoffs yields the lowest revenues of all envy-free equilibria. Both EOS and Varian identify the lowest-revenue envy-free (LREF) equilibrium as the most plausible. Varian offers an informal argument for the LREF equilibrium, while the main result of EOS shows that in a generalized English auction, the LREF equilibrium is the unique outcome. Cary et al. (2007) show that if bidders play myopic strategies in a repeated GSP auction, the system converges to the LREF equilibrium. A number of subsequent papers use the LREF equilibrium as the solution concept for sponsored search auctions.

However, it remains uncertain whether the LREF equilibrium is the "right" equilibrium to select in GSP auctions. Thompson and Leyton-Brown (2009) computationally explore the set of all equilibria in complete information GSP. Borgers et al. (2007) consider various Nash equilibria of the complete information game and suggest that a profit-maximizing search engine should be able to coordinate bidders on an equilibrium with revenue greater than LREF. However, our Proposition 4 shows that when a reserve price is chosen optimally, a search engine cannot coordinate bidders on an equilibrium with revenues higher than in the LREF equilibrium.

D. The Non-Contradiction Criteria (NCC)

As long as the outcome of the incomplete information game quickly converges to an equilibrium of the complete information game, one can view a complete information game as a valid approximation of the incomplete informa-

tion game. This suggests an equilibrium selection criteria: An equilibrium of complete information game can be ruled out if there does not exist an equilibrium of the corresponding game of incomplete information that converges to the same outcome. As a result, expected revenue in a "plausible" equilibrium of complete information game integrated over all possible realizations of bidder values cannot exceed the expected revenues in some equilibrium of the incomplete information game. We call this the Non-Contradiction Criteria (NCC).

Let us define NCC precisely. Denote by m the perfect Bayesian equilibrium of the incomplete information game that yields the largest expected revenue. Denote by $R_m(T)$ the expected revenue in equilibrium m in a world that lasts T periods. Let \bar{R} denote the (largest possible) expected per period equilibrium revenues in an incomplete information game $\bar{R} = \frac{1}{T} \lim_{T \rightarrow \infty} R_m(T)$. Consider equilibrium per-period revenues in an equilibrium μ of the complete information game. Denote by $r_\mu(\mathbf{s})$ the expected revenues in Nash equilibrium μ of a complete information GSP auction where values of bidders are given by vector \mathbf{s} .

DEFINITION 1: *A NE μ of a complete information game fails NCC if it generates greater expected revenues than any equilibrium of the corresponding incomplete information game. That is, μ fails NCC if $\int_{s_1} \dots \int_{s_N} r_\mu \mathbf{s} f(s_1) \dots f(s_N) ds_1 \dots ds_N > \bar{R}$.*

The idea behind NCC differs from the logic of most equilibrium refinements. Refinements such as subgame perfection, perfect Bayesian equilibria, and intuitive criteria result from ruling out strategies that are consistent with NE yet implausible in the actual game. In contrast, NCC calls for comparing *two distinct games*, say G and Γ . Γ is a game of complete information that is believed to approximate the behavior of players in a corresponding game of incomplete information \hat{G} . Some equilibria of Γ may correctly capture behavior in G while others may not. If we can establish some fact that is true in any equilibrium of G (e.g. the upper bound on equilibrium revenues), this fact can be used to narrow the set of relevant equilibria of Γ . Note that NCC is not a restriction on the set of equilibria in Γ ; NCC says nothing about reasonable

equilibria in Γ . Rather, NCC restricts equilibria of Γ when Γ is meant to capture behavior in G .

We use NCC to analyze sponsored search advertising. But the underlying idea is general. Many models of complete information games are simplified representations of complex dynamic games of incomplete information. In these cases, the set of plausible equilibria of a static game should be considered in the context of the corresponding incomplete information game. Our insight is that even if the incomplete information game is very complex, it may be possible to bound the set of equilibria of the incomplete information game.

E. Optimal Mechanism for the Incomplete Information Environment

Even though the sponsored search environment allows for multistate dynamic mechanisms, we will show that a one-shot generalized English auction is an optimal mechanism.

To apply NCC, we must find an upper bound on revenue in the dynamic mechanism. In this section, we describe the optimal mechanism for the incomplete information environment.

Consider the generalized English auction for the incomplete information environment. The auction is run once, and positions are allocated for the entire game. Following EOS, imagine a clock showing the current price, and continuously increasing over time. The price begins at the reserve price, and all advertisers willing to pay at least the reserve price participate in the auction. As the clock ticks upwards, advertisers can drop out. The auction is over when the next-to-last advertiser drops out. The last remaining advertiser is placed in the best (highest) position.

In general, optimal auction design for the sale of multiple heterogeneous objects remains an open problem. However, the structure of sponsored search makes it possible to extend the single-object results of Myerson (1981) to cover multiple items. Although the equilibrium price per click is higher in top positions, each advertiser's value for each position can be derived from a single variable – the advertiser's value per click. Ulku (2009) extends the Myerson framework to the case of multiple heterogeneous objects when bidders' private information

is one-dimensional.¹

Obtaining an optimal mechanism is a key intermediate result in our analysis. Similarly, Roughgarden and Sundararajan (2007) and Athey and Ellison (2009) also obtain an optimal mechanism as an intermediate result, but for entirely different purposes. Iyengar and Kumar (2006) consider computational methods to estimate optimal auctions when clickthrough rates are arbitrary. In contrast, the structure of clickthrough rates assumed in our model allows us to obtain an analytic result.

The following proposition characterizes the optimal mechanism²:

PROPOSITION 2: *The generalized English auction with reserve price r^* is an optimal mechanism with reserve price r^* that solves $r^* - \frac{1-F(r^*)}{f(r^*)} = 0$.*

PROOF:

The optimal direct revelation mechanism can be characterized using the same technique as in a single object case, except that the probability of receiving an object is replaced with the expected number of clicks that a bidder receives. Denote $x_k(\mathbf{r})$ the expected number of clicks received by bidder k when the realized vector of bidder values is given by $\mathbf{r} = (\mathbf{r}_1 \dots \mathbf{r}_N)$, and let $f(\mathbf{r})$ denote the pdf of vector \mathbf{r} and $t_k(0)$ the expected payment of bidder k when his value is zero. With this modification, formula 5.12 from Krishna (2002) gives seller revenues in a direct revelation mechanism: $\sum_{k=1}^N t_k(0) + \int_{\mathbf{r}} (\sum_{k=1}^N \psi(r_k) x_k(\mathbf{r})) f(\mathbf{r}) d\mathbf{r}$. Thus, only bidders with positive virtual valuations are allocated a positive expected number of clicks, and bidders with higher virtual (and actual) valuations are allocated higher positions. Thus, the generalized English auction achieves the same allocation and revenue as the optimal direct revelation mechanism.

COROLLARY 3: *The optimal reserve price in the generalized English auction is independent of*

¹Thus, even if bidders differ in exogenous quality, an optimal mechanism can be constructed.

²We thank Michael Ostrovsky for suggesting a dramatic simplification of the version of this proof that appeared in the 2006 draft of this paper.

the number of bidders, number of slots being auctioned, and the rate of decline in click-through rate from position to position.

F. Applying NCC

Consider application of NCC to a complete information GSP auction with reserve price r^* .

PROPOSITION 4: *In a complete information GSP auction with a reserve price of r^* , from among all envy-free equilibrium only the lowest revenue equilibrium survives NCC. Furthermore, this equilibrium is unique.*

PROOF:

According to EOS Theorem 1, all envy-free equilibria of the GSP auction yield revenues that are at least as large as revenue in the corresponding VCG auction. Yet Proposition 2 shows that these revenues equal the optimal revenue in the incomplete information environment. All but one envy-free equilibria of the complete information GSP auction yield higher revenues than \bar{R} (optimal revenue for incomplete information environment) and hence are ruled out by NCC.

II. Reserve Prices and Individual Bidder Payoffs

In this section, we explore the impact of reserve prices in an equilibrium of the generalized English auction where strategies are continuous in player types.

A. Bidders' Cost Increases

In this section, we will show that raising the reserve price creates an indirect effect far larger than the direct effect.

The reserve price *directly* affects the lowest bidder. As long as there are more positions than bidders, the lowest bidder pays the reserve price.) The reserve price may also *indirectly* affect other bidders because change in reserve price may impact equilibrium behavior of all players.

Consider an increase in a reserve price by an amount $\Delta r = r_{new} - r_{old}$. Suppose that there are more slots than advertisers, that n bidders had values greater than r_{old} , and that j bidders dropped out as a result of the increase in reserve.

In the following proposition, we compare the revenue increases resulting from a change in reserve price.

PROPOSITION 5: *If a changed reserve price leads j advertisers to drop out, then the total payment of each remaining advertiser changes by $\alpha_{n-j}\Delta r$.*

SKETCH OF PROOF:

For the full proof, see the online appendix.

Advertiser n receives slot n and pays reserve price r . Changing the reserve price by $\Delta r = r^1 - r^0$ changes n 's total payment $\alpha_n \Delta r$.

Advertiser $n - 1$ pays a per-click fee $p_{n-1} = b_n$. With reserve price r , advertiser n bids $b_n = s_n - \frac{\alpha_n}{\alpha_{n-1}}(s_n - r)$. (See Theorem 2 of EOS.) So the change in reserve price causes $n - 1$'s per-click payment to change by

$$\begin{aligned} \Delta p_{n-1} &= \Delta b_n = b_n^1 - b_n^0 \\ &= \left(s_n - \frac{\alpha_n}{\alpha_{n-1}}(s_n - r^1) \right) \\ &\quad - \left(s_n - \frac{\alpha_n}{\alpha_{n-1}}(s_n - r^0) \right) \\ &= \frac{\alpha_n}{\alpha_{n-1}} \Delta r \end{aligned}$$

Advertiser $n - 1$ receives α_{n-1} clicks, so his total payment changes by $\alpha_{n-1} \Delta p_{n-1} = \alpha_{n-1} \left(\frac{\alpha_n}{\alpha_{n-1}} \Delta r \right) = \alpha_n \Delta r$.

Recurring upwards yields an identical cancellation of terms at each step – hence the same total payment change, $\alpha_n \Delta r$, for each advertiser.

Now consider the case in which the changed reserve r^1 exceeds the valuations of the lowest j advertisers. Unable to achieve a positive profit, these advertisers exit. Advertiser $n - j$ takes the role of n in the preceding analysis, and each advertiser's total payment changes by $\alpha_{n-j} \Delta r$.

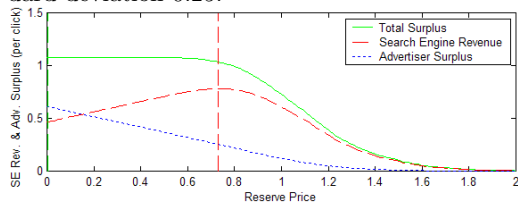
COROLLARY 6: *If no advertisers drop out as a result of a reserve price change, then the reserve price yields an identical dollar-for-dollar change in the total payment of each advertiser.*

COROLLARY 7: *In a market with n bidders bidding above the reserve price, a reduction in reserve price has an indirect effect at least $n - 1$ times greater than the direct effect. (We say that an advertiser is affected indirectly by a reduction in reserve price if his payment before the reduction was strictly greater than the reserve price.)*

B. Welfare Implications

Under realistic assumptions, an optimal reserve price can yield a notable increase in search engine revenue. For example, in October 2008, Yahoo CEO Sue Decker stated in Yahoo’s Q3 2008 earnings call that Yahoo’s new reserve price efforts were the “most significant” cause of the 11% increase in Yahoo’s revenue per search.

Through simulations, we have explored how reserve price impacts both bidders and the seller. The plot below shows results of simulations with five advertisers whose values are drawn from a log-normal distribution with mean 1 and standard deviation 0.25.



The graph reveals that moving from a zero reserve price to optimal reserve price leads to a small decline in total surplus and a significant increase in search engine revenues.

III. Conclusion

In online search advertising auctions, reserve prices are important – offering large revenue benefits to sellers, along with large welfare effects for the market as a whole. But the effect of reserve price can be opaque, particularly because it is not always straightforward to compare “before” and “after” conditions. Our Proposition 5 assesses the distributional consequences, revealing the counterintuitive across-the-board effects of reserve price increases that might have seemed to target low bidders exclusively.

Our paper also offers insight on equilibrium selection. When one game is used to understand results of another, careful comparison can yield useful insights. This approach applies in many contexts where a one-shot game offers a simplified view of repeated interactions, e.g. Cournot and Bertrand models of industry competition.

We see this paper as an example of *applied theory*. We address the important applied problem of setting an optimal reserve price in a large and growing market, and our approach analyzes

mechanisms that are widely used. Yet our theoretical contribution – equilibrium selection criteria – also grounds our results. Moreover, our counterintuitive revenue finding – an indirect effect larger than the direct effect – stands at the intersection of theory and application.

REFERENCES

- Z. Abrams and M. Schwarz. Ad auction design and user experience. *Applied Economics Research Bulletin*, pages 98–105, October 2008.
- S. Athey and G. Ellison. Position auctions with consumer search. NBER Working Paper No. 15253, 2009.
- T. Borgers, I. Cox, M. Pesendorfer, and V. Petricek. Equilibrium bids in sponsored search auctions: Theory and evidence. University of Michigan Ann Arbor Manuscript, 2007.
- M. Cary, A. Das, B. Edelman, I. Giotis, K. Heimerl, A. Karlin, C. Mathieu, and M. Schwarz. Greedy bidding strategies for keyword auctions. *Proceedings of the ACM Conference on Electronic Commerce*, 2007.
- B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, pages 242–259, April 2007.
- G. Iyengar and A. Kumar. Characterizing optimal adword auctions. CORC Report 2006-04.
- V. Krishna. *Auction Theory*. Academic Press, 2002.
- R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- T. Roughgarden and M. Sundararajan. Is efficiency expensive? In *Third Workshop on Sponsored Search Auctions, WWW2007*.
- D. Thompson and K. Leyton-Brown. Computational analysis of perfect-information position auctions. In *Proceedings of the ACM Conference on Electronic Commerce*, 2009.
- Levent Ulku. Optimal combinatorial mechanism design. Centro de Investigacion Economica, ITAM, Working Paper 0903, 2009.
- H. Varian. Position auctions. *International Journal of Industrial Organization*, pages 1163–1178, October 2007.

IV. Online Appendix: Proof of Proposition 5

PROPOSITION 5: If a changed reserve price leads j advertisers to drop out, then the total payment of each remaining advertiser changes by $\alpha_{n-j}\Delta r$.

PROOF:

Setup. Advertisers are $1..n$, and slots are also $1..n$. Initial reserve price r^0 changes to some new reserve price r^1 . Suppose for now that the new reserve price r^1 remains sufficiently low that no advertiser is priced out of the market, *i.e.* $r^1 \leq s_n$.

Advertiser n . Advertiser n receives slot n and pays the reserve price r . So when the reserve price increases from r^0 to r^1 , advertiser n 's per-click payment increases by $\Delta r = r^1 - r^0$. Advertiser n receives α_n clicks, so advertiser n 's total payment increases by $\alpha_n \Delta r$.

Advertiser $n-1$. The per-click fee p_{n-1} paid by advertiser $n-1$ is determined by the per-click bid b_n of advertiser n . With reserve price r , advertiser n bids $b_n = s_n - \frac{\alpha_n}{\alpha_{n-1}}(s_n - r)$. (This is the base case of Theorem 2 of EOS.) So when the reserve price changes from r^0 to r^1 , advertiser $n-1$'s change in per-click payment equals advertiser n 's change in per-click bid, which is

$$\begin{aligned} \Delta p_{n-1} &= \Delta b_n = b_n^1 - b_n^0 \\ &= \left(s_n - \frac{\alpha_n}{\alpha_{n-1}}(s_n - r^1) \right) \\ &\quad - \left(s_n - \frac{\alpha_n}{\alpha_{n-1}}(s_n - r^0) \right) \\ &= \frac{\alpha_n}{\alpha_{n-1}} \Delta r \end{aligned}$$

Advertiser $n-1$ receives α_{n-1} clicks. So when the reserve price increases by Δr , advertiser $n-1$'s total payment increases by $\alpha_{n-1} \Delta p_{n-1} = \alpha_{n-1} \left(\frac{\alpha_n}{\alpha_{n-1}} \Delta r \right) = \alpha_n \Delta r$.

Advertiser $n-2$ and the general case. Advertiser $n-2$ pays a per-click fee p_{n-2} given by b_{n-1} . With reserve price r , advertiser $n-1$ bids $b_{n-1} = s_{n-1} - \frac{\alpha_{n-1}}{\alpha_{n-2}}(s_{n-1} - b_n)$. (This is the general case of Theorem 2 of EOS.) So when the reserve price changes from r^0 to r^1 , advertiser

$n-1$'s change in bid is

$$\begin{aligned} \Delta b_{n-1} &= b_{n-1}^1 - b_{n-1}^0 \\ &= \left(s_n - \frac{\alpha_{n-1}}{\alpha_{n-2}}(s_k - b_n^1) \right) \\ &\quad - \left(s_n - \frac{\alpha_{n-1}}{\alpha_{n-2}}(s_k - b_n^0) \right) \\ &= \frac{\alpha_{n-1}}{\alpha_{n-2}} (b_n^1 - b_n^0) \\ &= \frac{\alpha_{n-1}}{\alpha_{n-2}} (\Delta b_n) \\ &= \frac{\alpha_{n-1}}{\alpha_{n-2}} \left(\frac{\alpha_n}{\alpha_{n-1}} (r^1 - r^0) \right) \\ &= \frac{\alpha_n}{\alpha_{n-2}} \Delta r \end{aligned}$$

Advertiser $n-2$ receives α_{n-2} clicks. So advertiser $n-2$'s total payment increases by $\alpha_{n-2} \Delta b_{n-1} = \alpha_{n-2} \frac{\alpha_n}{\alpha_{n-2}} \Delta r = \alpha_n \Delta r$. Recursing upwards to advertiser 1 confirms that each advertiser's total payment increases by the same amount, $\alpha_n \Delta r$.

Reserve price that excludes one or more advertisers. We now allow an increase in reserve price such that one or more advertisers is priced out of the market (relaxing the assumption in the first paragraph of the proof).

Suppose the increased reserve r^1 exceeds advertiser n 's valuation s_n . Then advertiser n can never achieve a positive profit by buying ads; n would have to pay more than his valuation. So n exits. What about the other advertisers? Advertiser $n-1$ now takes on the role of n in the preceding analysis, and $n-1$'s payment per click increases one-for-one with the reserve price (above the point where n dropped out). Then $n-2$ takes the role of $n-1$ in the preceding analysis, and similarly recursing upwards. If one advertiser, n , drops out, then each advertiser's total payment increases by $\alpha_{n-1} \Delta r$. More generally, if the increased reserve price leads j advertisers to drop out, then each other advertiser's total payment increases by $\alpha_{n-j} \Delta r$.