ECAI 2010
H. Coelho et al. (Eds.)
IOS Press, 2010
© 2010 The authors and IOS Press. All rights reserved.
doi:10.3233/978-1-60750-606-5-587

Designing a Successful Adaptive Agent for TAC Ad Auction

Meng Chang and **Minghua He**[1] and **Xudong Luo**[2]

Abstract. This paper describes the design and evaluation of AstonTAC, the runner-up in the Ad Auction Game of 2009 International Trading Agent Competition. In particular, we focus on how AstonTAC generates adaptive bid prices according to the Market-based Value Per Click and how it selects a set of keyword queries to bid on to maximise the expected profit under limited conversion capacity. Through evaluation experiments, we show that AstonTAC performs well and stably not only in the competition but also across a broad range of environments.

1 Introduction

In recent years, sponsored search [3, 7] has become the indispensable source of revenue for Internet search engine companies like Google, Yahoo and MSN/Bing. Instead of showing the same advertisement to every user, it enables companies to promote their products to targeted groups of consumers based on their search queries [8]. Moreover, pricing advertisement is through keyword auction which has a number of advantages over the conventional negotiation between the seller and the buyer such as price efficiency and resource-saving. Thus, sponsored search has become one of the most efficient and profitable forms of advertising and attracted considerable research attention from various fields.

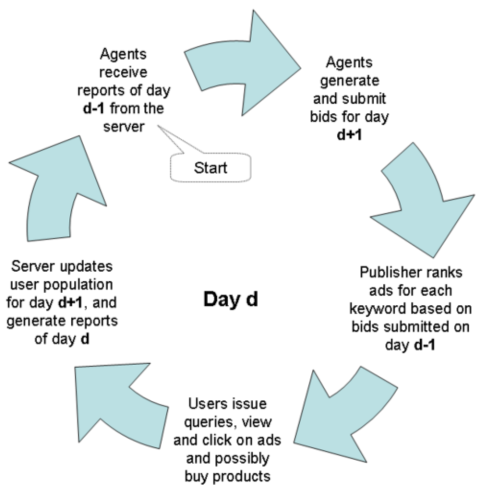

Figure 1. Activities cycle of each day in TAC AA 2009

The investigation of sponsored search generally falls into three categories: (i) search user behaviour modelling, (ii) mechanism design to improve its effectiveness and efficiency, (iii) strategy formulation faced mainly by advertisers in order to maximise their return from a budget-constrained advertising campaign. The Trading Agent Competition Ad Auction (TAC AA[3]) has provided an ideal test bed for advertiser strategies. In its scenario, there are three kinds of agents: advertisers, publishers and users. The behaviour of the publishers and users are generated by TAC AA server according to some fixed stochastic policy [5]. There are eight advertiser agents (entrants to the competition) that compete against each other on ad placement for search result given queries over 60 days of bidding periods.

In more detail, agents represent retailers of home entertainment products featured by three manufacturers (*Flat, Lioneer* and *PG*) and three components (*TV, Audio* and *DVD*). Altogether, there are nine distinct products. A query generated by a user is a (*manufacturer, component*) pair and unspecified manufacturer or component is denoted as '$null$'. In total, there are 16 possible queries at three focus levels denoted by $F0, F1, F2$. The more specific the query, the higher the focus level.

$$\begin{aligned}
Q_{F0}:\ & \{(null, null)\}\\
Q_{F1}:\ & \{(Flat, null), (Lioneer, null), (PG, null), (null, TV),\\
& (null, Audio), (null, DVD)\}\\
Q_{F2}:\ & \{(Flat, TV), (Flat, Audio), (Flat, DVD),\\
& (Lioneer, TV), (Lioneer, Audio), (Lioneer, DVD),\\
& (PG, TV), (PG, Audio), (PG, DVD)\}
\end{aligned}$$

On each day of the game and for each query type above, an auction is run to determine the ad placements. Once an ad is clicked and leads to a customer transaction, it is called a conversion. Based on the focus levels, distributions of click probability P_{click}, continuation probability $P_{continuation}$ (probability that a user proceeds to click the next ad), and conversion rate $P_{conversion}$ differ between queries. In addition, each advertiser is assigned a Manufacturer Speciality (*MS*) and a Component Speciality (*CS*) in each game. If a query matches *MS*, it will receive a high conversion value. If a query matches *CS*, it will receive a high conversion rate. The number of conversion is softly constrained by *Distribution Capacity* C^{cap}. TAC AA introduces C^{cap} to impose the effect of diminishing marginal value [5] of conversion: when the number of conversion increases to some point, $P_{conversion}$ will start to drop and result in lower conversion profit due to increased cost.

On each day of the game and for each query type, the advertiser agent submits a bid to the publisher. Such a bid specifies the bid price (the maximum amount that an advertiser is willing to pay for a click on his ad), the spend limit (the corresponding ad will be excluded once the spend limit is reached) and the ad display type (either

[1] Aston University,United Kingdom, email: {changm2, m.he1}@aston.ac.uk
[2] Nanyang Technological University, Singapore, email: xdluo@ntu.edu.sg

[3] http://aa.tradingagents.org

generic or *targeted*). At the end of the game, agents are evaluated based on their cumulative surplus: sales profit less cost they paid for all the clicks received in the game.

Figure 1 shows activities of each day of the game. On day *d*, each advertiser agent first receives market reports of day *d-1*, then decides a bid for each query to submit for day *d+1*. The publisher ranks bids submitted by different advertiser agents on day *d-1* and works out a price per click to charge on each query for each agent. As users click on ads and buy products from advertiser agents, the server collects every agent's impression, position, clicks, conversion, revenue and cost to generate market reports of day *d*. Bidding period begins from day *0* and agents' first bid submission is for day *1*. Consequently market reports are available from day *2*.

According to specific settings of TAC AA, designing a successful agent mainly faces two challenges: **a)** without knowledge of bid prices of other participants, what is the optimal position and how to decide an appropriate bid price to target it? **b)** given non-deterministic conversion limit, how to maximise profit in terms of both number of conversion and conversion value? To address the first problem, we find an alternative way to build bid prices on Market-based Value Per Click (MVPC) which can be estimated based on system parameters or market reports. To overcome the second challenge, we estimate the true maximum number of conversion allowed by distribution capacity and select only the most profitable queries to bid on and fill the expected conversion allowance every day.

The remainder of this paper is organised as follows. Section 2 describes our agent. Section 3 analyses its performance and further evaluate it by controlled experiments. Finally, Section 4 concludes.

2 AstonTAC

AstonTAC is composed of four components: Agent Knowledge Base, Bid Price Generator, Query Selector and Ad Display Selector shown in Figure 2.

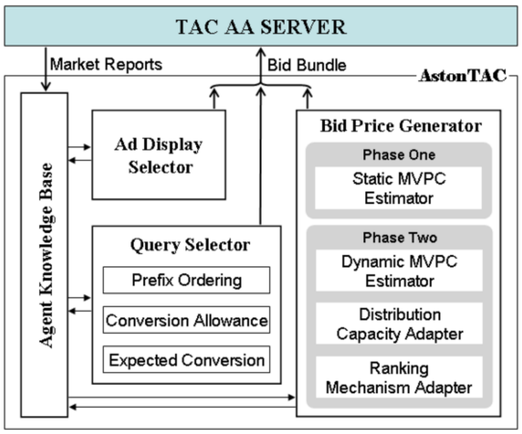

Figure 2. Architecture of AstonTAC

Agent Knowledge Base is designed to process, organise and record information from the server and turn them into knowledge for other components to use. The information it deals with can be divided into two categories: static information including setting parameters and game initialisation information and dynamic run-time information such as market reports received on daily basis. The other three components collectively generate the bid $B=\langle b, l, t\rangle$ for each query q and submitted as a bundle on each day of the game. In particular, Bid Price Generator (Section 2.1) calculates a bid price b. Query Selector (Section 2.2) specifies a spend limit l. Ad Display Selector (Section 2.3) makes the choice of ad display type t.

2.1 Bid Price Generator

Based on the availability of market reports, the whole game can be divided into two phases: Phase One (day *0* and *1*) and Phase Two (from day *2* to *59*). Bid price generation in both phases is based on the same concept - Market-based Value Per Click (MVPC) - of each query.

Definition 1 A query's MVPC is the expected conversion revenue minus advertising cost that a click on its ad can generate.

We introduce MVPC because we find clicks a kind of special commodity as they are not retainable. In an auction of normal commodity, the bidder is willing to pay up to his/her valuation for an item because he/she can own the item after the payment. However, in a keyword auction, clicks cannot be owned so that bidders always want to make a profit out of them. For this reason, we believe MVPC is the true worth of a click to an advertiser. Unlike the conventional Value Per Click [2, 9], MVPC is much more dynamic by incorporating the cost. Basically, if it is assumed that revenue-per-click is independent of position [6] which is disapproved by [4], MVPC still varies as cost-per-click (CPC) is expected to vary for different positions. Moreover, due to conversion limit, one query's conversion affects the others' conversion rates rendering the change of revenue-per-click such that MVPC reflects the real value of a query in an interdependent multi-query environment. Therefore, if we can estimate every query's MVPC and set a bid price accordingly, our bids should automatically approximate the best response to the market without explicitly targeting any position. MVPC is estimated in two different ways for different phases of the game.

2.1.1 Phase One

In phase one, MVPC is estimated using expected revenue-per-click multiplied by a fixed discount ratio $r_{discount}$ indicating the proportion of profit in revenue. Expected revenue-per-click is a product of expected conversion value and conversion rate. Because MVPC is based on static fundamental information, it is denoted as v_{static},

$$v_{static} = P_{conversion} \cdot v_{con} \cdot r_{discount} \tag{1}$$

The conversion value v_{con} in above formula is given by,

$$v_{con} = \begin{cases} USP(1+MSB) & \text{if } q_m = \textit{MS} \\ \frac{2}{3}USP + \frac{1}{3}USP(1+MSB) & \text{if } q_m = \textit{null} \\ USP & \text{otherwise} \end{cases} \tag{2}$$

where USP (Unit Sale Price) is the standard conversion value, MSB=50% is *MS* bonus rate [5] and q_m denotes the manufacturer part of a query. The first branch of the above formula means q_m=*MS* queries receive a 50% higher conversion value than the normal USP. The second branch means q_m=*null* queries has 1/3 probability to receive a higher conversion value and 2/3 probability to receive a normal conversion value. The third branch means if a query's manufacturer part is a specific manufacturer other than *MS*, its conversion value will be exactly USP.

$P_{conversion}$ is calculated in a similar way,

$$P_{conversion} = \begin{cases} \frac{(1+CSB)\pi_l}{1+CSB\pi_l} & \text{if } q_c = \text{CS} \\ \frac{2}{3}\pi_l + \frac{1}{3}\frac{(1+CSB)\pi_l}{1+CSB\pi_l} & \text{if } q_c = \text{null} \\ \pi_l & \text{otherwise} \end{cases} \tag{3}$$

where q_c denotes the component part of a query, CSB=50% is *CS* bonus rate and π_l (l$\in$ $\{F0, F1, F2\}$) denotes the baseline conversion rate. A query's conversion rate is primarily confined by its π_l in accordance with its focus level. On this basis, if q_c=*CS*, a superior conversion rate calculated by function $\eta(x,y) = \frac{xy}{xy+(1-x)}$ where x=π_l and y=$(1+CSB)$ will apply. If a query's q_c is *null*, it has 1/3 chance to receive the superior conversion rate because generally 1/3 user population's preference matches the advertiser's *CS*. If a query's q_c is a specific component other than the advertiser's *CS*, its conversion rate will be exactly its corresponding π_l.

The discount ratio $r_{discount}$ in Formula 1 is a value between 0 and 1. The best $r_{discount}$ is chosen so that AstonTAC's accumulative profit in first five day of the game is maximised in experimental games.

By now, v_{static} has been calculated. Apparently, queries with relatively high v_{static} are expected to produce more profit on each conversion and vice versa. Since slightly higher bid price does not necessarily raise the cost due to the features of Generalised Second Price [2, 3] mechanism, it is sensible to set high bid prices for relatively high-value queries to increase their chance of receiving top positions. To counterbalance the possible increase of conversion from high-value queries under restricted total conversion, relatively low bid prices should be set for relatively low-value queries. To this end, h_q ($80\% \leq h_q \leq 120\%$) is introduced to enforce the heuristic. In addition, higher C^{cap} means more conversions and clicks are acceptable. Therefore, we also introduce $h_c \in \{85\%, 100\%, 115\%\}$[4] so that AstonTAC generally bids lower when $C^{cap} = 300$ or higher when $C^{cap} = 500$. Together bid of each query in phase one is given by,

$$b_0 = v_{static} \cdot h_q \cdot h_c \tag{4}$$

2.1.2 Phase Two

In phase two, dynamic MVPC denoted by $v_{dynamic}$ is calculated according to dynamic market reports. Consequently, a query's bid b in phase two is built on its $v_{dynamic}$ and adjusted by distribution capacity adapter δ and ranking mechanism adapter β simultaneously.

$$b = v_{dynamic} \cdot \delta \cdot \beta \tag{5}$$

Dynamic MVPC Formula 6 incorporates information about a query's revenue, cost and clicks from recent W (W is the size of aggregation window for distribution capacity. Specifically, W=5) days to calculate average profit of a click so far as the expected profit a click can possibly make on the next day.

$$v_{dynamic} = \frac{\sum_{i=d-1}^{d-W} revenue_i - \sum_{i=d-1}^{d-W} cost_i}{\sum_{i=d-1}^{d-W} click_i} \tag{6}$$

We aggregate W-day data for two reasons: (i) C^{cap} takes effect on the basis of W-day aggregate number of conversion. (ii) it reduces the unwanted fluctuation of $v_{dynamic}$ caused by system dynamics. As a result of the above formula, $v_{dynamic}$ is highly responsive and adaptive to the change of environment represented by three key factors:

1. *Search user population* - it determines the baseline number of possible clicks, conversion and revenue
2. *Competition intensity* - it determines the ad rank and cost-per-click
3. *Conversion probability* - it determines the number of clicks needed to generate a conversion

The rationale behind is that we map the above three factors to the three components *revenue, cost and clicks* respectively in Formula 6. *revenue* is an indicator of user population - a larger number of active users tend to generate more revenue assuming other two factors are fixed. *Cost* is an indicator of competition intensity - assuming user population is fixed and C^{cap} is not filled, the severer the competition, the more it costs to maintain the same position. At last, *number of clicks* is an indicator of the conversion probability - excess conversion beyond C^{cap} will cause lower $P_{conversion}$ which means more clicks are needed to generate the same number of conversions.

Distribution Capacity Adapter Distribution Capacity Adapter is designed to adapt our bid prices to $C^{cap} \in \{300, 400, 500\}$ a decisive factor in TAC AA explicitly. It strongly confines the advertiser's potential profit by affecting conversion rates. According to the focus levels, each agent's query is set a default $P_{conversion,def}$ by the server. During the game, once W-day accumulative conversion exceeds C^{cap}, the timely conversion rate $P_{conversion,t}$ will start to drop below $P_{conversion,def}$. When $P_{conversion,t}$ is sufficiently low, clicks will make losses rather than profits because users only click on the ad but almost never make a transaction. Hence, there is a trade-off between number of conversions and conversion rates at which a critical number of conversion $C_{crit}(C_{crit} > C^{cap})$ occurred such that profit of next conversion equals to its cost. The aim of setting a bidding constraint is to keep them both high so that accumulative profit can be maximised. C^{cap} is a soft constraint because exceeding C^{cap} does not stop conversion but only reduce its probability. C_{crit} can be estimated (see Section 2.2) but cannot be used here because bids are only allowed to be changed daily rather than every time a conversion happens. Eventually, the ratio δ between C^{cap} and expected W-day aggregate conversion by day *d* $c_{agg,d}$ is found to be a suitable indicator of bid adjustment over $v_{dynamic}$. To obtain δ, $c_{agg,d}$ is first estimated using weighted average,

$$c_{agg,d} = \frac{\sum_{t=d-1}^{0} (w_t \sum_{i=t}^{t-(W-1)} c_i)}{\sum_{t=d-1}^{0} w_t} \tag{7}$$

where exponential weight $w_t = \omega^{d-t-1} \quad (0 < \omega < 1)$ and c_i denotes total conversion from all queries on day *i*. Then the adjustment factor δ is given by,

$$\delta = \frac{C^{cap}}{c_{agg,d}} \tag{8}$$

The intuition behind δ is: if $\delta > 1$, C^{cap} is expected to be underfilled on day *d*, then all bids are increased by δ, then the number of clicks is expected to increase as well as the subsequent conversions on day *d+1*; if $\delta < 1$, all bids are reduced for the opposite effect. In this game, δ falls in the range between 0.6 and 1.8.

Ranking Mechanism Adapter Ranking mechanism adapter adjusts the bid price further by taking into account of the ranking mechanism adopted by the publisher. TAC AA employs a squashing parameter $\chi(0 \leq \chi^4 \leq 1)$ initialised at the beginning of each game to interpolate between two extremes: $\chi = 0$ is equivalent to rank-by-bid and $\chi = 1$ is equivalent to rank-by-revenue [7]. Specifically, given e_q as the estimated click through probability by the publisher for query q and b_q as the bid on q, the ranking score is calculated as $b_q(e_q)^\chi$. In order to adapt our bid prices to the dynamic ranking mechanism ranking mechanism adapter β is introduced and unknown e_q (e_q employed by the publisher is not revealed to the advertiser) estimated

[4] h_c, h_q and $r_{discount}$ are chosen by the domain expert based on experiment results.

using the aggregate Click-Through-Rate (CTR) and denoted as e'_q,

$$\beta = \left(1 + e'_q\right)^{-\chi} = \left(1 + \frac{\sum_{i=0}^{d-1} Click_q}{\sum_{i=0}^{d-1} Impression_q}\right)^{-\chi} \quad (9)$$

Based on β, bid price b_q stays unchanged if $e'_q = 0$, or $\chi = 0$. Otherwise, b_q is reduced with the increase of either e'_q or χ.

2.2 Query Selector

Given limited conversion capacity, query selector selects only a set of queries to bid on such that the expected available conversions are allocated to the queries that can potentially generate high profit. The following algorithm shows the selection process.

1. Prefix-ordering queries.
2. Estimate conversion allowance $C_{w,d+1}$ for day *d+1*.
3. Estimate expected conversion of each query $c_{q,d+1}$ on day *d+1*.
4. Identify a bidding set of queries $A = \{1, \ldots, s\}$, $s \in \{1, 16\}$:
    - 4.1. IF $C_{w,d+1} \leq 0$ THEN $A = \phi$; GOTO 5. ELSE initialise s to 1.
    - 4.2. WHILE $\sum_{q \in A} c_{q,d+1} < C_{w,d+1}$ and $s < 16$ DO $s = s + 1$.
    - 4.3. $A = \{1, 2, \ldots, s\}$.
5. Set $l = \infty$ for each $q \in A$ and $l = 0$ for each $q \notin A$.

As we can see, the output of query selector is a set of spend limits. Selected queries are not restricted by a spend limit, thus their spend limit is set to infinite. Unselected queries will not be active on day *d+1* so that their spend limit is set to zero. The intuition here is: by estimating conversion allowance, the maximum number of conversion before expected conversion profit drops to zero is revealed; by prefix-ordering and selecting first s queries to put in the bidding set A, we make sure the allowance is used by high-profit queries only. The remainder of this section explains how Step 1, 2 and 3 work.

Prefix-ordering Prefix-ordering [8] is used to sort and prefix queries in the descending order of their profit-per-conversion (PPC). A query's PPC is calculated as follows,

$$PPC = \frac{\sum_{i=d-1}^{0} revenue_i - \sum_{i=d-1}^{0} cost_i}{\sum_{i=d-1}^{0} conversion_i} \quad (10)$$

where values of variables are obtained from market reports.

Conversion Allowance Referring to the discussion of distribution capacity adapter in Section 2.1.2, expected conversion allowance is the difference between C_{crit} and conversions of four recent days including day d which can be estimated as $c_{agg,d} - c_{d-4}$, ($c_{agg,d}$ is given by Formula 7). Once C^{cap} is exceeded, every additional conversion lowers timely conversion rate $P_{conversion,t}$ by $\lambda = 0.995$. Hence, we model C_{crit} as $C_{crit} = C^{cap} + n$ where n is the number of additional conversions by which $P_{conversion,t}$ reaches a critical value $P_{conversion,crit}$ such that expected conversion revenue equals to conversion cost (clicks needed to generate a conversion times CPC). At this equilibrium point, n is maximised. Since both conversion revenue and CPC differ across queries, we introduce general conversion revenue v'_{con} as the average revenue with respect to total conversion and general CPC c'_{click} as average cost with respect to total clicks from all queries,

$$v'_{con} = \frac{\sum_{i=d-1}^{0} \sum_{q \in all} revenue_{q,i}}{\sum_{i=d-1}^{0} \sum_{q \in all} conversion_{q,i}} \quad (11)$$

$$c'_{click} = \frac{\sum_{i=d-1}^{0} \sum_{q \in all} cost_{q,i}}{\sum_{i=d-1}^{0} \sum_{q \in all} click_{q,i}} \quad (12)$$

Mathematically, the equilibrium point can be presented as,

$$v'_{con} = P^{-1}_{conversion,crit} \cdot c'_{click} \quad (13)$$

Since $P_{conversion,crit}$ equals to $P_{conversion,std} \times \lambda^n$, we have

$$n = \log_\lambda \frac{P_{conversion,crit}}{P_{conversion,std}} \quad (14)$$

where $P_{conversion,std}$ is an average of baseline conversion rates weighted by the distribution of both queries and search population towards different focus levels. n decreases with c'_{click} because the larger the conversion cost the less excess conversion is needed to reach the equilibrium point. n increases with v'_{con} because the larger the conversion revenue, it takes more excess conversion to bring $P_{conversion,t}$ down from $P_{conversion,std}$ to $P_{conversion,crit}$. Once n is found, conversion allowance $C_{w,d+1}$ can also be obtained,

$$C_{w,d+1} = C^{cap} + n - (c_{agg,d} - c_{d-4}) \quad (15)$$

Expected Conversion We model expected conversion of each query on day *d+1* as a product of expected impression, click probability and conversion rate,

$$C_{q,d+1} = Impression_{d+1} \cdot P_{click,d+1} \cdot P_{conversion,d+1} \quad (16)$$

$Impression_{d+1}$ is estimated based on impressions occurred on all queries in last $Pr^{-1}_{burst} = 10$ days where Pr_{burst} is the search population burst rate. $P_{conversion,d+1}$ is estimated as a product of $P_{conversion}$ given by Formula 3 and $\sqrt{\delta}$ (δ is given by Formula 8),

$$P_{conversion,d+1} = P_{conversion} \min\left(1, \sqrt{\delta}\right) \quad (17)$$

$P_{click,d+1}$ is dependent on the relevant bid which is already generated by bid price generator and stored in knowledge base. We first estimate an exponential function [6] for each query to map bid to position. Then we infer $P_{click,d+1}$ according to distributions of click probability and continuation probability provided in the game specification of TAC AA 2009.

2.3 Ad Display Selector

Finally we discuss how to choose an ad display type t between *Generic* and *Targeted*. Generic ad leads to query's system default click-through-rate whereas targeted one can either brings the effective click-through-rate over or under the system default one depending on whether query's component part matches user's underlying component preference. Our following heuristic rule works well in the competition.

$$t = \begin{cases} Generic & \text{if } q_c \neq \mathit{CS} \text{ and } q_m \neq \mathit{MS} \\ Targeted & \text{if } q_c = \mathit{CS} \text{ or } q_m = \mathit{MS} \end{cases} \quad (18)$$

Based on this rule, for (non-MS,CS) queries or (MS,non-CS) queries, targeted ad will cause a lower-than-default click probability from the users whose underlying product preference disagrees with our product speciality. However, this is not a truly adverse result. First, users with another component preference are less likely to buy our specialised component. It is pointless to display a generic ad which increases the odds of clicks leading to more cost. Secondly, if a user with different manufacturer preference purchases a product made by our specialised manufacturer, as the manufacturers do not match, our profit is only the standard value. Because the underlying query is a MS query, we expect more clicks coming from users with the same preference to purchase *MS* products and yield a larger profit of USP(1+MSB). Since distribution capacity is limited, for users with different manufacturer preference, we would rather like them less likely to click on our ads such that the chance of their low-revenue-same-cost conversion becomes smaller.

3 Evaluation

In this section, AstonTAC is analysed from two aspects: competition results to identify the successful properties and three controlled experiments to test robustness of our agent.

3.1 Game Results and Analysis

In TAC AA 2009, AstonTAC ranked 2nd out of 15 teams in both qualifying games and the final. In the final, 40 games were played on server one and server two simultaneously and respectively. We download logs of all forty games run on server one for analysis in which we are particularly interested in the top three agents - TacTex, AstonTAC and Schlemazl - whose average scores in the final are \$79886, \$76281 and \$75408, respectively.

We start with a correlation test to see whether agent's profit potential is affected by C^{cap}. It turns out correlation coefficient between C^{cap} and average profit is over 97.7%. This proves the importance of adapting bidding strategies to C^{cap}. Such strong correlation also implies that it is only appropriate to make comparison of performance between agents based on same setting of C^{cap} or identical number of each different C^{cap} in case of analysing aggregate results.

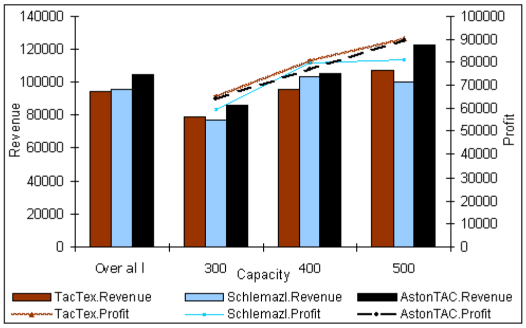

Figure 3. Average revenue and profit of top three agents. The left vertical axis shows the revenue and the right vertical axis shows the profit. The four groups of bars, from left to right, represent the averaged revenue of agents with different capacity: overall (all games), 300, 400 and 500 respectively. The three curves show the correlation between the profit and capacity for each of the three agents.

As can be seen in Figure 3, AstonTAC performs the best in terms of revenue generation. We believe there are two reasons. First, we set high bid prices for high-value queries to target top positions. High-value means high expected profit per click and top position brings maximum number of clicks. Second, we suppress low-profit conversions by bidding less on low-value queries and selecting only profitable keywords to bid on. Moreover, for AstonTAC and TacTex, profit forms a clear ascending trend against capacity whereas for Schlemazl there no significant profit increase from $C^{cap} = 400$ to $C^{cap} = 500$. The particularly low profit and revenue at $C^{cap} = 500$ indicates that Schlemazl did not sufficiently exploit its conversion space in high capacity. With 25% incremental capacity brought by the change of C^{cap} from 400 to 500, Schlemazl's average number of conversion only increases by 0.96%. By contrast, TacTex's increase rate is 12.3% and AstonTAC's is 14.23%.

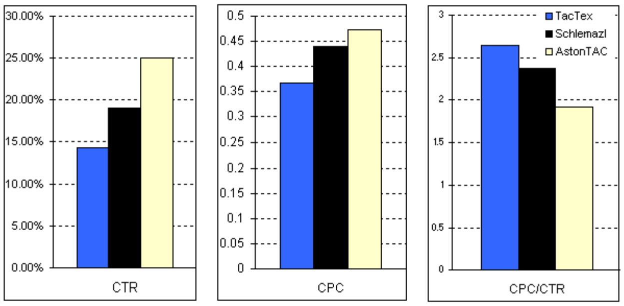

(a) (b) (c)

Figure 4. (a) Average CTR (b) Average CPC (c) CPC/CTR of AstonTAC, Schlemazl and Tactex.

CTR is one of the most important criteria to judge whether an online advertising campaign is successful. Furthermore, in rank-by-revenue mechanism, high historical CTR can enhance rank and lower payment for the same bid. In TAC AA, high CTR is particularly preferred because parameter χ is always closer to 1 than 0 meaning that the ranking mechanism is more by-revenue than by-bid. Based on this, we set high bid prices for high-value queries to target the 1st slot and highest possible CTR. Conversely, our high CTR results in comparatively low CPC. Figure 4 shows although our average cost per click is larger than the other two, but our advantage in CTR justifies it. The smallest CPC/CTR means our cost increasing speed against CTR is slower than the other two. To sum up, the rise of cost is dominated by the rise of revenue, our strategy benefits more from increased revenue than suffers from increased cost.

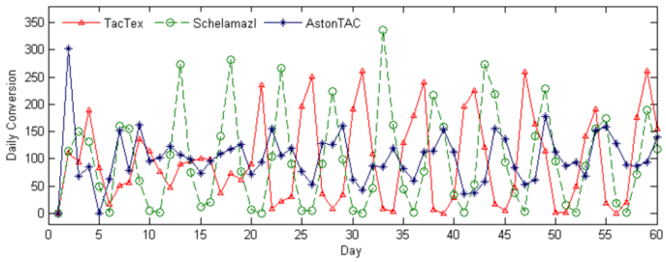

Figure 5. Different variance of top three agents in daily conversion.

Moreover, AstonTAC performs well not because it makes sky-high profit on some days but because it consistently makes large profit while other agents could earn nothing periodically. This can be justified by the low variance in daily conversion. Figure 5 demonstrates this through visualizing daily conversions of AstonTAC, TacTex and Schlemazl in a typical final game where C^{cap} is the same for each agent. As Figure 5 shows, with statistically identical mean, AstonTAC's standard deviation is only 55.9% and 51.2% of that of TacTex and Shlemazl, respectively. This should largely be attributed to behaviour associated with C_{crit} which stetches our ability of making conversions even when C^{cap} is exceeded.

3.2 Controlled Experiments

In order to see whether AstonTAC works well in a broader range of environment such as competing with other agents than participants in TAC AA, we have purposefully designed three controlled experiments. Table 1 shows settings of each experiment.

Table 1. Controlled game settings

Exp.	Participants	Capacity Setting	Rounds
A	AstonTAC, TacTex, Dummy $\times 6$	Game default	80
B	AstonTAC, TacTex, AstonBB, AstonRules, Dummy $\times 4$	Game default	80
C	AstonTAC, TacTex, Dummy $\times 6$	Identical C^{cap}	15

Experiment A and B are set in the same way of TAC AA. What is changed is the participating agents. In Experiment C, all agents are assigned identical $C^{cap} \in \{300, 400, 500\}$ in each game.

3.2.1 *Experiment A*

AstonTAC is the overall winner in Experiment A. In fact, out of 80 games, AstonTAC won 54 whereas TacTex only won the rest 26. AstonTAC's average score is $50263(\pm 7944.6)$ whereas TacTex got $45439(\pm 8254.9)$. Besides, the deterioration of return-on-investment and cost-per-click comparing with competition results for TacTex are more than that of AstonTAC. We believe our relatively stable performance is in connection with the unpredictable environment caused by dummy agents who are expected to exercise stochastic bidding. Because our bids are based on query's value, AstonTAC is less affected by environmental unpredictability than other agents. In this experiment, both AstonTAC and TacTex's overall performances are worse than that in the final suggesting that social welfare can only be achieved when every agent bids wisely.

3.2.2 *Experiment B*

Two more agent are introduced in Experiment B: AstonBB and AstonRules. AstonBB was initially developed essentially based on balanced bidding strategy [1]. AstonRules employs heuristic rules to infer bid according to the position each query receives in each round. AstonTAC is the overall winner again but with a very small margin over TacTex. However, our strategy seems quite superior for high capacity as AstonTAC's average profit at C_{500} is \$3185 more than TacTex. We believe both *Distribution Capacity Adapter* and *Query Selector* contribute significantly to this result. As for query selector, the algorithm works better at high C^{cap} because prefix decision gets preciser with larger capacity. As the number of intelligent agent increases in the game, TacTex's performance tends to increase rapidly whereas AstonTAC still being stable. It implies that TacTex may have the ability to recognise the bidding pattern of other intelligent agents and act accordingly to undercut the intelligence of their strategy. In contrast, AstonTAC's strategy is holistically built on the basis of dynamic market-based value per click, which does not need to target any specific position such that it cannot be easily undercut. For this reason, it appears to present stable and reliable performance in whatever environment especially unpredictable ones.

3.2.3 *Experiment C*

In this experiment, agents' performance can be compared directly because capacity bias is eliminated. Out of a total of fifteen games, AstonTAC has won ten (2/3). AstonTAC's average profit is \$50412 and TacTex's average profit is \$47269. Figure 6 shows the overall CTR of TacTex and AstonTAC with the change of capacity (300, 400 and 500). It can be seen that TacTex's click-through-rate forms a declining trend against capacity while AstonTAC's trend is ascending which is more compatible with intuitions. Our attention to the most crucial game parameter C^{cap} and adaptive bidding strategy designed through different components accordingly should be the explanation. For the same reason, it is a significant feature for AstonTAC that its profit increases with the capacity.

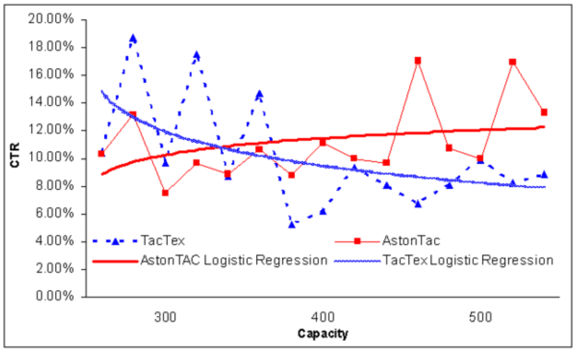

Figure 6. CTR trend against change of distribution capacity.

4 Conclusion

AstonTAC has been shown to be successful and stable across a wide range of TAC AA environments both in the competition and in controlled experiments. In particular, we attribute the success of AstonTAC to the strategy used by the bid price generator and the query selector. Market-based Value Per Click reflects the dynamic change of the market and thus leads to the generation of flexible and adaptive bidding prices.

The strategies employed here are tailored to the specific context of the AA competition. However, due to the similar features of the TAC ad competition and real sponsored search, we believe that concepts developed for AstonTAC are broadly applicable to an advertiser agent in a real sponsored search scenario. One of our future work is to integrate our MVPC-based strategy with large-scale budget-constraint optimisation approach [6] and apply the agent in a real pay-per-click auction.

REFERENCES

[1] M. Cary, A. Das, B. Edelman, I. Giotis, K. Heimerl, and A. R. Karlin, 'Greedy bidding strategies for keyword auctions', *In ACM EC*, (2007).

[2] B. Edelman, M. Ostrovsky, and M. Schwarz, 'Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords', *Workshop on Sponsored Search Auctions*, (2006).

[3] D. Fain and J. Pedersen, 'Sponsored search: A brief history', *Bulletin of the American Society for Information Science and Technology 32(2)*, 12–13, (2006).

[4] A. Ghose and S. Yang, 'An empirical analysis of search engine advertising: Sponsored search in electronic markets', *Management Science Vol. 55, No. 10, pp. 1605-1622*, (2009).

[5] P. R. Jordan, B. Cassell, et al., 'The ad auctions game for the 2009 trading agent competition', *Computer Science & Engineering University of Michigan Ann Arbor, MI 48109-2121 USA*, (2009).

[6] B. Kitts and B. Leblanc, 'Optimal bidding on keyword auctions', *Electronic Markets*, (2004).

[7] S. Lahaie and D. M. Pennock, 'Revenue analysis of a family of ranking rules for keyword auctions', *Eighth ACM Conference on Electronic Commerce, San Diego*, 50–56, (2007).

[8] P. Rusmevichientong and D. Williamson, 'An adaptive algorithm for selecting profitable keywords for search-based advertising services', *Proceedings of the 7th ACM conference on electronic commerce*, (2006).

[9] Y. Vorobeychik, 'Simulation-based game theoretic analysis of keyword auctions with low-dimensional bidding strategies', *Conference on Uncertainty in Artificial Intelligence*, (2009).