Ex-post IR Dynamic Auctions with Cost-per-action Payments*

Extended Abstract

Weiran Shen
IIIS, Tsinghua University
Beijing, China
emersonswr@gmail.com

Zihe Wang
ITCS, Shanghai University of Finance and Economics
Shanghai, China
wang.zihe@mail.shufe.edu.cn

Song Zuo
IIIS, Tsinghua University
Beijing, China
songzuo.z@gmail.com

ABSTRACT

Consider a repeated auction between one seller and many buyers, where each buyer only has an estimation of her value in each period until she actually receives the item in that period. The seller is allowed to conduct a dynamic auction to sell the items but must guarantee ex-post individual rationality. In other words, if the buyer realized that her value of the item she just received was zero, she did not need to pay anything. Unlike the clicks on the ads, these actions are private information only observable by the buyers (advertisers). Hence they may have incentives to misreport the user actions, because they can pay less under cost-per-action payment schemes with ex-post individual rationality guarantees.

In this paper, we use a structure that we call *credit accounts* to enable a general reduction from any *incentive compatible* and *ex-ante individual rational* dynamic auction to an *approximate incentive compatible* and *ex-post individually rational* dynamic auction with credit accounts. Our reduction can obtain stronger individual rationality guarantees at of the cost of weaker incentive compatibility. Surprisingly, our reduction works without making any common knowledge assumptions. Finally, as a complement to our reduction, we prove that there is no non-trivial auction that is exactly incentive compatible and ex-post individually rational under this setting.

CCS CONCEPTS

• **Theory of computation → Algorithmic game theory and mechanism design**; **Computational advertising theory**; *Computational pricing and auctions*; • **Applied computing → Online auctions**;

KEYWORDS

Dynamic auctions; ex-post individual rationality; cost-per-action payments; credit accounts; ad auctions

ACM Reference Format:
Weiran Shen, Zihe Wang, and Song Zuo. 2018. Ex-post IR Dynamic Auctions with Cost-per-action Payments. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

*The authors thank the anonymous reviewers for their helpful comments. The authors benefit from discussions with Balasubramanian Sivan and Renato Paes Leme. This paper is supported in part by the National Natural Science Foundation of China Grant 61561146398, a China Youth 1000-talent Program, an Alibaba Innovative Research Program, the Shanghai Sailing Program (Grant No. 18YF1407900) and the Fundamental Research Funds for the Central Universities.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

1 INTRODUCTION

Internet advertising has been playing a very important role in the advertising industry. Most online advertising platforms, such as search engines and social media, have gone through the evolution from the *cost-per-mille impressions* (CPM) model to the *cost-per-click* (CPC) model, where the former is aligned with traditional advertising while the latter focuses more on performance. In the CPC model, when a user requests a certain web page, the platform collects bids from the advertisers and based on these bids, determines whose advertisement to display on the page. The corresponding advertiser is charged when her advertisement is clicked by the user. Such an advertising model is called the CPC model because the advertiser only needs to pay when her advertisement is clicked. This CPC model has been the de facto model for most major online advertising platforms, and is proven to be profitable [13]. However, despite its success, this model is criticized to have the click fraud problem, i.e., the competitors of an advertiser, or even the platform itself, may deliberately create false clicks to increase the advertiser's cost or to extract more revenue. Furthermore, the advertisers have to pay for clicks that do not lead to final purchase of their products. Although one may argue that in expectation the advertisers are indeed profitable, it may still be a serious problem for small companies that cannot ignore such risks.

A relatively new model that has gained more research attention recently is the *cost-per-action* (CPA) advertising model. In contrast to the CPC model, the CPA model is even more performance-oriented and focuses directly on user actions on the advertiser's web page. In the CPA model, the advertisers are only charged when the users make certain actions, such as purchases or transactions. It seems that the CPA model and the CPC model are almost the same except for the payment. However, this advertising model clears the uncertainty faced by the advertiser and can potentially decrease the vulnerability to click fraud. Besides these advantages, the CPA model also gives more incentives to the platforms to deliver high-quality impressions to the users. In 2007, the CPA model was described as the "Holy Grail" of targeted advertising by Google [29]. Currently, many online advertising platforms, including Google, eBay, Amazon, Facebook, Baidu and WeChat have already started to test the CPA model.

Another essential difference between these two models is that the platform cannot directly observe the users' actions on the advertisers' websites whereas the users' clicks are observable by both the platform and the advertiser. Such an undesirable property may cause the advertisers to hide the users' actions to avoid payments. This also poses challenges in putting the CPA model in practice to replace the CPC model that is currently dominant in the advertising industry.

This paper is directly motivated by the above challenge. In this paper, we aim to tackle the incentive problem and present a new auction mechanism called the *credit account mechanism*. Our mechanism solves the incentive issue by setting a credit account for each advertiser and follows the "allocate-report-pay" scheme. In our mechanism, the advertisers are given a certain amount of "credit quota" and an advertiser cannot win the auction if her credit runs out of her "quota". Once an impression is allocated after a periodic auction, the advertiser reports back to the platform her value of the action taken by the user. The mechanism then charges this advertiser by some amount less than the reported value and updates the advertiser's credit by the difference between the price she actually paid and the expected per-impression payment. Intuitively, such a credit account works as a tolerance for hiding user actions, since an advertiser's credit quickly runs out in that case. However, an honest advertiser only has a negligible chance of consuming all her credit.

Our contributions. The contributions in this paper are briefly summarized as follows:

- We formalize a framework that we call credit accounts.[1] Using this framework, we can reduce any general incentive compatible and ex-ante individually rational mechanism to a credit account mechanism that can implement the same allocation rule as the original mechanism with high probability and guarantees approximate incentive compatibility and ex-post individual rationality.
- Such a reduction naturally induces a trade-off between the strength of the approximate incentive compatibility and the probability of desired implementation. In particular, it also applies to second price auctions.
- As a complement to the constructed credit account mechanisms, we show that the exact incentive compatibility and the ex-post individual rationality cannot be achieved simultaneously, unless the mechanism is trivial. In this sense, credit account mechanisms have achieved the strongest properties we can hope for.

1.1 Related Works

Ever since Myerson's seminal paper on designing revenue optimal auctions [26], there have been intensive researches on analyzing and designing one-shot auction mechanisms. For example, Edelman et al. [13] and Varian [32] study the performance of the generalized second price auction (GSP), which is the mostly widely used mechanism among major search engines in the world. Hartline and Roughgarden [16], Shen and Tang [28] and Bachrach et al. [6] provide mechanisms that can tradeoff among different objectives (e.g. revenue, welfare and click yields). There is also a rich literature on multi-item auctions [2, 10–12, 19, 30, 33–35], and on repeated auctions motivated by online advertising [3, 4, 7, 14, 18, 31].

[1] In fact, we are note the first to use the idea of credit accounts in mechanism design. Similar reputation based structures are used for other settings [15, 20].

A closely related line of work is dynamic mechanism design (see Bergemann and Välimäki [9] for a comprehensive survey). For example, Athey and Segal [5] provide an efficient, budget-balanced and Bayesian incentive compatible mechanism in the dynamic setting. Bergemann and Välimäki [8] consider repeated auctions of a single item, where all buyers' values are independent. They focus on efficient allocations and give a mechanism called the dynamic pivot mechanism, which is similar to the second price auction. Mierendorff [21] studies a dynamic setting where each buyer has a deadline for buying the item. They give sufficient conditions such that the deadline constraints can be fulfilled or violated. He also gives the optimal auction mechanism when there are two buyers and two periods.

There is also a series of works that focus on designing mechanisms with the CPA advertising model. Nazerzadeh et al. [27] study the setting where the advertisers' value may evolve over time. They present a mechanism that satisfies asymptotic individual rationality and asymptotic incentive compatibility. However, their mechanism does not exactly fall into the CPA advertising model, since the winner still needs to pay even if the user does not click on his advertisement. Hu et al. [17] compare the CPC advertising model and the CPA advertising model. Their results show that the CPA model is better in incentivizing the platform to improve the purchase rate, but suffers from the adverse selection problem. Agarwal et al. [1] consider a similar setting where the advertisers report both the predefined actions and the action probabilities. They show that at equilibrium, the advertisers may report skewed bids. However, their results only hold in one-shot games.

Our proposed mechanism also benefits from some highlevel ideas of the "bank account" mechanism, where the seller maintains a "bank account" for each buyer during the dynamic auction [22–25]. Although with similar names, the "credit account" in this paper is fundamentally different from the "bank account": (i) the bank account mechanisms are designed under the common knowledge assumption to ensure dynamic incentive compatibility, while the credit account mechanism guarantees approximate dynamic incentive compatibility *without* any common knowledge assumption; (ii) the "balance" in bank accounts can be thought of as money, where the buyers might be charged through their bank accounts, while the "credit" in the credit accounts is more like a "score" that measures the reliability of the buyers based on their past behaviors.

REFERENCES

[1] Nikhil Agarwal, Susan Athey, and David Yang. 2009. Skewed bidding in pay-per-action auctions for online advertising. *The American Economic Review* 99, 2 (2009), 441–447.
[2] Saeed Alaei, Hu Fu, Nima Haghpanah, Jason Hartline, and Azarakhsh Malekian. 2012. Bayesian optimal auctions via multi-to single-agent reduction. In *Proceedings of the 13th ACM Conference on Electronic Commerce.* ACM, 17–17.
[3] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. 2013. Learning prices for repeated auctions with strategic buyers. In *Advances in Neural Information Processing Systems.* 1169–1177.
[4] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. 2014. Repeated contextual auctions with strategic buyers. In *Advances in Neural Information Processing Systems.* 622–630.
[5] Susan Athey and Ilya Segal. 2013. An efficient dynamic mechanism. *Econometrica* 81, 6 (2013), 2463–2485.
[6] Yoram Bachrach, Sofia Ceppi, Ian A Kash, Peter Key, and David Kurokawa. 2014. Optimising trade-offs among stakeholders in ad auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation.* ACM, 75–92.
[7] Santiago Balseiro, Max Lin, Vahab Mirrokni, Renato Leme, and Song Zuo. 2017. Dynamic Revenue Sharing. In *Advances in Neural Information Processing Systems.*

2681–2689.
[8] Dirk Bergemann and Juuso Välimäki. 2010. The dynamic pivot mechanism. *Econometrica* 78, 2 (2010), 771–789.
[9] Dirk Bergemann and Juuso Välimäki. 2017. Dynamic Mechanism Design: An Introduction. (2017).
[10] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. 2012. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM, 459–478.
[11] Constantinos Daskalakis. 2015. Multi-item auctions defying intuition? *ACM SIGecom Exchanges* 14, 1 (2015), 41–75.
[12] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2013. Mechanism design via optimal transport. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 269–286.
[13] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review* 97, 1 (2007), 242–259.
[14] Alessandro Epasto, Mohammad Mahdian, Vahab Mirrokni, and Song Zuo. 2018. Incentive-Aware Learning for Large Markets. In *Proceedings of the 27th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee.
[15] Chen Hajaj, John P Dickerson, Avinatan Hassidim, Tuomas Sandholm, and David Sarne. 2015. Strategy-Proof and Efficient Kidney Exchange Using a Credit Mechanism.. In *AAAI*. 921–928.
[16] Jason D Hartline and Tim Roughgarden. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, 225–234.
[17] Yu Hu, Jiwoong Shin, and Zhulei Tang. 2015. Incentive problems in performance-based online advertising pricing: cost per click vs. cost per action. *Management Science* 62, 7 (2015), 2022–2038.
[18] Yash Kanoria and Hamid Nazerzadeh. 2014. Dynamic reserve prices for repeated auctions: Learning from bids. (2014).
[19] Xinye Li and Andrew Chi-Chih Yao. 2013. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences* 110, 28 (2013), 11232–11237.
[20] Yuan Liu, Jie Zhang, Han Yu, and Chunyan Miao. 2014. Reputation-Aware Continuous Double Auction.. In *AAAI*. 3126–3127.
[21] Konrad Mierendorff. 2016. Optimal dynamic mechanism design with deadlines. *Journal of Economic Theory* 161 (2016), 190–222.
[22] Vahab Mirrokni, Renato Paes Leme, Rita Ren, and Song Zuo. 2018. Dynamic Mechanism Design in the Field. In *Proceedings of the 27th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee.
[23] Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. 2016. Dynamic auctions with bank accounts. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
[24] Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. 2016. Optimal dynamic mechanisms with ex-post IR via bank accounts. *arXiv preprint arXiv:1605.08840* (2016).
[25] Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. 2018. Non-Clairvoyant Dynamic Mechanism Design. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM.
[26] Roger B Myerson. 1981. Optimal auction design. *Mathematics of operations research* 6, 1 (1981), 58–73.
[27] Hamid Nazerzadeh, Amin Saberi, and Rakesh Vohra. 2013. Dynamic pay-per-action mechanisms and applications to online advertising. *Operations Research* 61, 1 (2013), 98–111.
[28] Weiran Shen and Pingzhong Tang. 2017. Practical versus optimal mechanisms. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 78–86.
[29] Stephan Spencer. 2007. Google deems cost-per-action as the 'Holy Grail'. CNET News. (2007).
[30] Pingzhong Tang and Zihe Wang. 2017. Optimal mechanisms with simple menus. *Journal of Mathematical Economics* 69 (2017), 54–70.
[31] Pingzhong Tang and Yulong Zeng. 2018. The Price of Prior Dependence in Auctions. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM.
[32] Hal R Varian. 2007. Position auctions. *international Journal of industrial Organization* 25, 6 (2007), 1163–1178.
[33] Zihe Wang and Pingzhong Tang. 2014. Optimal mechanisms with simple menus. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 227–240.
[34] Andrew Chi-Chih Yao. 2014. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 92–109.
[35] Andrew Chi-Chih Yao. 2017. Dominant-strategy versus bayesian multi-item auctions: Maximum revenue determination and comparison. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 3–20.