Games and Economic Behavior 102 (2017) 583–623

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Computational analysis of perfect-information position auctions

David R.M. Thompson, Kevin Leyton-Brown *

Computer Science Department, University of British Columbia, Canada

A R T I C L E I N F O

Article history:
Received 20 May 2014
Available online 4 March 2017

JEL classification:
C63
C88
D44
D47
L86
M37

Keywords:
Computational mechanism analysis
Position auctions
Sponsored search
Compact game representations

A B S T R A C T

After experimentation with other designs, major search engines converged on weighted, generalized second-price auctions (wGSPs) for selling keyword advertisements. Theoretical analysis is still not able to settle the question of why they found this design preferable to other alternatives. We approach this question in a new way, adopting an analytical paradigm we dub "computational mechanism analysis." Specifically, we sample position auction games from a given distribution, encode them in a computationally efficient representation language, compute their Nash equilibria, and calculate economic quantities of interest. We considered seven widely studied valuation models from the literature and three position auction variants. We found that wGSP consistently showed the best ads of any position auction, measured both by social welfare and expected number of clicks. In contrast, we found that revenue was extremely variable across auction mechanisms and was highly sensitive to equilibrium selection, the preference model, and the valuation distribution.

1. Introduction

Position auctions are a relatively new family of mechanisms in which bidders place a single bid for a set of goods of varying quality, and the ith-highest bidder wins the ith-most-desirable good. Each year, these auctions yield billions of dollars selling advertising space on search-engine results pages. Various position auction designs have been considered over the years; e.g., advertisers who attract more clicks may be given advantages over weaker bidders, and bidders may have to pay their own bid amounts or a smaller amount computed from the bids of others. After some initial experimentation with these design dimensions, the major search engines have converged on a single design: the weighted, generalized second-price auction (which we denote wGSP; we define it formally in what follows). The main question that this paper seeks to address is whether wGSP represents a good choice, as compared both to the auctions it has replaced and to theoretical benchmarks. Specifically, we ask whether wGSP is more economically efficient, whether it generates more revenue, whether it yields results that users find more relevant, and whether it produces low-envy allocations.

There is an enormous literature on auction analysis that seeks to answer such questions. Overwhelmingly, this literature proceeds by modeling a setting as a (Bayesian; perfect-information) game and then using theoretical analysis to describe what occurs in (Bayes–Nash; dominant-strategy; locally envy-free) equilibrium. There is much to like about this approach: it is often capable of determining that a given mechanism optimizes an objective function of interest or proving that no

* Corresponding author.
E-mail addresses: daveth@cs.ubc.ca (D.R.M. Thompson), kevinlb@cs.ubc.ca (K. Leyton-Brown).

http://dx.doi.org/10.1016/j.geb.2017.02.009

mechanism can satisfy a given set of properties. Indeed, most of what we know about mechanism design was established through such analysis. However, the approach also has limitations: in order to obtain clean theoretical results it can be necessary to make strong assumptions about bidder preferences and to simplify tie-breaking and bid discretization rules. Even when considering such a simplified version of a given problem, it is often extremely difficult to make quantitative comparisons between non-optimal mechanisms (e.g., which non-revenue-optimizing mechanism yields higher revenue on expectation)? The current state of affairs in the literature is thus that we know a great deal about auctions, but that many open questions appear to be resistant to analysis by known techniques.

In the case of position auctions, most research has used perfect-information Nash equilibrium as the solution concept of choice. This choice has been justified by the fact that advertisers interact repeatedly: nearly identical goods—user views of a particular search result page—are sold up to millions of times per day, and advertisers can continually adjust their bids and observe the effects.[1] A variety of other technical assumptions are commonly made, characterizing advertiser preferences (e.g., a click has the same value regardless of position, and regardless of which other ads are shown), user behavior (e.g., a user's response to an ad is independent of which other ads she has seen), and advertiser behavior (e.g., an advertiser will act to reduce his envy even when doing so does not increase his utility). With all these assumptions in place, various strong results have been obtained (e.g., wGSP is efficient and generates weakly more revenue than VCG), but other important questions remain open (e.g., does wGSP generate more revenue than other position auctions?). Relaxing any one of these assumptions can lead to many further open questions.

This paper introduces *computational* techniques for the analysis of mechanisms and shows that these techniques are able to address a wide variety of open questions that have not proven amenable to theoretical analysis. (For a survey of related methods, see Section 2.3.) We maintain the approach of modeling a mechanism as a game and of reasoning about (exact) solution concepts of interest; we depart from traditional analysis by allowing only a discrete set of bids and by answering questions by providing statistical evidence rather than theorems. Specifically, we sample advertiser preferences from a given distribution, compute an exact Nash equilibrium of the resulting game between advertisers, and reason about this equilibrium to compute properties of the outcome, such as expected revenue or social welfare. By repeatedly sampling, we can make quantitative statistical claims about our position auction setting (e.g., that one auction design generates significantly more expected revenue than another). While our approach is not specific to position auctions, we focus on that domain here to demonstrate that our methods are able to yield qualitatively new findings about a widely studied setting.

This approach, which we dub "computational mechanism analysis", differs in many ways from theoretical methods. As already mentioned, one clear disadvantage is that our approach only produces statistical results (e.g., given distribution D, A performs significantly better than B on expectation) rather than theorems (e.g., A always performs better than B). Conversely, our methods have the advantage that they are able to produce results in settings for which such simple patterns do not exist. For example, we have observed distributions over advertiser preferences under which the wGSP auction sometimes generates far less revenue than its predecessor, uGSP, and sometimes generates far more. Thus, we know that any comparison of these auctions must necessarily be statistical and distribution dependent, rather than guaranteeing that one of these auctions always yields more revenue. Our computational approach also allows us to consider arbitrary preference distributions, possibly derived from real-world data, as opposed to being restricted to distributions with convenient theoretical properties like monotone hazard rates. A further property of our computational mechanism analysis approach is both a benefit and a weakness: bidders must be restricted to a finite set of discrete bids, unlike the vast majority of literature on auction theory, which assumes that bids can take any value in a real interval. While we depart from this tradition, we do not see discreteness as necessarily disadvantageous; real-world position auctions tend to be rather coarsely discrete. For example, position auctions often clear for tens of cents per click, while bids are required to be placed in integer numbers of cents. Finally, some auction features, like rules for tie-breaking and rounding, are difficult to analyze in the continuous case, but pose no obstacle to our approach.

Given all of these advantages, the reader might wonder why the computational analysis of mechanisms is not already commonplace. The answer is that to date, the scale of the problem has made computational equilibrium analysis infeasible for virtually all games of interesting size. Normal-form games grow exponentially in the number of players, making all but the simplest problems too big even to store on the most powerful of modern computers. Even when representation is not a problem, finding a sample Nash equilibrium of a general-sum game is computationally hard, with every known algorithm requiring time that grows exponentially in the size of the normal form. These problems compound, meaning that overall, equilibrium computation requires time doubly exponential in the number of players! Luckily, recent advances in algorithmic game theory (surveyed in Section 2) offer a way around this roadblock: games can be encoded in exponentially less space and equilibria can be identified exponentially more quickly when players' payoffs are suitably structured. In Section 3 we describe algorithms for encoding many position auction variants into such efficient representations, guaranteeing exponential improvements in both representation size and runtime. Previously, it was infeasible to compute exact Nash equilibria of position auction games even using a supercomputer. Using our approach, exact Nash equilibria can often be found on a desktop PC in less than a second.

[1] Despite these strong arguments in favor, we note that it is possible to disagree with the perfect-information model: even under repeated interaction, bidders never explicitly observe each others' values; what they do observe may be obfuscated by search engines; and the setting changes as bidders enter and exit the market. Theoretical analysis shows that the imperfect information model can yield qualitatively different outcomes (Gomes and Sweeney, 2014).