

Simplified Mechanisms with an Application to Sponsored-Search Auctions

Paul Milgrom¹

First draft: August 16, 2007

This draft: December 17, 2008

A mechanism can be simplified by restricting its message space. If the restricted message spaces satisfy a certain “outcome closure property,” then the simplification is “tight”: for every $\varepsilon \geq 0$, any ε -Nash equilibrium of the simplified mechanism is also an ε -Nash equilibrium of the unrestricted mechanism. Prominent auction and matching mechanisms are tight simplifications of mechanisms studied in economic theory and often incorporate price-adjustment features that facilitate simplification. The generalized second price auction used for sponsored-search advertising is a tight simplification of a series of second-price auctions that eliminates the lowest revenue equilibrium outcomes and leaves intact only higher revenue equilibria.

Keywords: sponsored search, generalized second-price auctions, mechanism design.

JEL Categories: D44, C78

I. Introduction

Real-world resource allocation mechanisms are often much simpler than the direct mechanisms featured in economic theory.² The reason is not hard to see: when there are

¹ Support for this research was provided by National Science Foundation Grant SES-0648293 and by Yahoo! Thanks to Marissa Beck and Richard Steinberg for comments and to Joshua Thurston-Milgrom for editorial help. Any opinions expressed here are those of the author alone.

² A *direct mechanism* is a mechanism in which each participant’s strategy is a message describing that participant’s private information, or “type.”

more than a few distinct items to be allocated, the sheer number of combinations of items makes it too costly for participants to determine and express all of their relevant values, as direct mechanisms ordinarily require.

The resource allocation mechanisms used in practice often employ messages that are too simple to describe preferences completely. For example, in simultaneous first-price auctions of the sort utilized for wholesale trading of used cars to dealers, the auctioneer typically accepts individual bids on cars, and allows the bidder no opportunity to describe the extent to which it might be willing to substitute one car for another. Similarly, the National Resident Matching Program (NRMP) uses a variant of the celebrated Gale-Shapley algorithm³ to assign doctors to hospitals, but accepts reports from hospitals that consist only of a number of positions and a rank order list of doctors, allowing a hospital only a meager opportunity to describe its preferences about the composition of its incoming class.

For our theoretical analysis, we view the short messages in these examples as reports of preferences from a restricted set. The bids on cars in our example describe the values of a hypothetical bidder with additive preferences, and the NRMP messages describe the relevant choices by a hospital which, given a choice between two doctors, would make that choice independently of the other doctors that it has hired. In our theory, a *simplified mechanism* is derived from an original mechanism by restricting the set of messages.

How does such a simplification affect the performance of the mechanisms observed in practice? This question is problematic, because a simplified mechanism may

³ Gale and Shapley (1962).

correspond to more than one *extension*. Different extensions assign different outcomes to the excluded message profiles, and any of these might be compared to the outcome of the simplification. Still, there are specific extensions that are of obvious interest.

One is the menu auction of Bernheim and Whinston (1986), which is an extension of the first-price auction that allows bidders to express separate values or bids for every package. The auctioneer accepts the collection of bids that maximizes the total value of the allocation and charges each bidder the price it has bid for its assigned package. For a full preference report, this mechanism entails a number of bids that grows exponentially with the number of lots or items sold. How much is a well-informed bidder harmed by a simplification that restricts messages to ones that describe additively separable preferences when other bidders are similarly restricted? The answer? None at all. If the bidder would win a collection of lots S at a price of p in the menu auction, then it could win the same collection with additive bids for the individual items that sum to no more than p an arbitrarily small amount. Of course, for a bidder who was less confident about its opportunities, the restriction to additive bids could be quite costly.

A similar analysis applies to the National Resident Matching Program. The NRMP algorithm selects the *doctor-best stable matching*⁴ for the preferences reported by doctors and hospitals. This selection is no accident; it was an explicit goal of the mechanism's design.⁵ Although the current NRMP algorithm allows hospitals to report

⁴ A matching of doctors to hospitals is *stable* if three conditions are satisfied: (1) no hospital prefers to dismiss any of its doctors, (2) no doctor prefers to quit his hospital, and (3) no doctor-hospital pair can make a deal that both prefer (taking into account that, after the deal, the hospital may choose to dismiss some of its assigned doctors). A *doctor-best stable matching* is a stable matching M with the property that there is no doctor D and stable matching M' such that D strictly prefers his hospital in M' to his hospital in M .

⁵ Roth and Peranson (1999).

only a limited set of preferences, it would be easy to extend it to allow hospitals to report preferences from any class for which doctor-best stable matchings always exist. The largest such class that also includes all the preferences that are currently reportable is the class of substitutable preferences.⁶

One may ask: how much harm does a particular hospital suffer when its messages are restricted to those allowed by the NRMP, rather than to the wider class of substitutable preferences? To investigate that question, fix a hospital H , reports by the doctors, and simple reports by the other hospitals. Is hospital H hurt by its inability to use some substitutable preference report? If hospital H is sufficiently well informed, then for any set of doctors S that hospital H could hire by making the alternative report, it can hire the same set by reporting that it has $|S|$ openings and that the doctors in S are at the top of its list. For the well-informed hospital, there is no harm at all, paralleling our finding for the sealed-bid package auction.

These examples demonstrate that some actual mechanisms can be regarded as simplifications of interesting extensions in a way that satisfies an *outcome-closure property*, which we define informally as follows: for every participant n and every profile of restricted messages m_{-n} and every unrestricted message \hat{m}_n for participant n , there is some restricted message m_n for n such that the message profiles (m_n, m_{-n}) and (\hat{m}_n, m_{-n}) both lead to the same allocation for n .⁷ Besides being satisfied by our sample mechanisms, this outcome-closure property is interesting because it represents a limited

⁶ The existence of a doctor-best stable match for substitutable preferences follows from results of Kelso and Crawford (1982) and is discussed by Roth and Sotomayor (1990). The converse is proven by Hatfield and Milgrom (2005).

⁷ The formal definition will allow approximate equality of outcomes rather than requiring exact equality.

but intuitively understandable notion of the richness of the message space and because, as shown below, it implies a limit on the set of possible pure Nash equilibria of the simplified direct mechanism.

Simplifying a mechanism by restricting the messages can change the set of equilibria in two ways. It can eliminate a Nash equilibrium profile by eliminating one of the messages in that profile. But a simplification can also introduce an additional Nash equilibrium by eliminating all the profitable deviations from some message profile. For some applications, it will be important to evaluate not only Nash equilibria but also ε -Nash equilibria, which are strategy profiles from which no player can gain more than ε by a deviation.⁸ We will say that a simplification is *tight* if for all admissible preferences over outcomes and for every $\varepsilon \geq 0$, every pure-strategy ε -Nash equilibrium of the simplified direct mechanism is also an ε -Nash equilibrium of the original mechanism.

According to our Simplification Theorem, any simplification that satisfies the outcome closure property is tight. Thus, regardless of hospitals' actual preferences concerning class composition, every pure equilibrium of the actual NRMP mechanism is an equilibrium of the extended version of the mechanism. Similarly, regardless of bidders' actual preferences regarding their bundles of used cars, every equilibrium of the item-by-item used car auction is an equilibrium of the corresponding menu auction.

A *direct auction mechanism* is a direct mechanism in which outcomes are goods allocations and cash payments. A nearly universal simplification used by direct auction mechanisms is *conflation*: participant's messages are required to report the same value

⁸ Such ε -Nash equilibria are fundamental to the study of "large" markets, in which participants expect that they can have only a small effect on prices and cannot gain much by misreporting their preferences. See Roberts and Postlewaite (1976).

for all potential allocations that assign the same bundle of goods to that bidder, with no opportunity to vary its value for that bundle based on others' assignments and payments. The Vickrey auction is a conflation of the Clarke-Groves mechanism, because the Vickrey auction has a narrower message space, but uses its messages to compute allocations and payments by just the same formula as is used for the Clarke-Groves mechanism.⁹

In commodities markets, conflation extends much further. Although many commodities differ on the basis of continuous variables such as size, color, location of origin, water content, etc, they are usually graded into a mere handful of discrete categories for purposes of pricing. Electricity pricing provides a similar example. In northern California, power delivered by producers at any of the approximately two hundred grid connection points are priced on the basis of just three regional prices corresponding to the California-Oregon border, the southern-northern transition point (SP-15), and the northern California injection point (NP-15).

Because any failure to distinguish among different products can lead to pricing errors, conflation is often accompanied by price adjustments to mitigate these errors. One example is the auction of rough diamonds regularly conducted by BHP Billiton. The company first sorts its diamonds into a small number of lots that are chosen to be as uniform as possible. It then sells "slices" of each lot by auction. A bidder may bid for a random slice from a lot, but may not specify any particular slice. The price paid for a particular slice is set by multiplying the auction price by an adjustment factor to account for the seller's estimate of variations in value among slices.

⁹ Vickrey (1961), Clarke (1971), Groves (1973).

A related example is the design chosen by the spectrum regulators in India and the UK for their sales of radio spectrum in the 2.6 GHz range, which are expected to occur in 2009. In those auctions, bidding will first determine the quantities of spectrum allocated to each bidder and set base prices to be paid by the bidders. There will then follow an assignment stage of bidding, in which winning bidders, who are already guaranteed to acquire spectrum, bid additional amounts to obtain particular collections of spectrum bands.¹⁰ These spectrum auction designs are hybrids, combining a conflation step with an assignment step that accounts for differences in spectrum quality.

Conflation is also an important design element in the world's most frequently used auction – the sponsored-search auction employed by many Internet search engines, including Google, Yahoo!, and MSN. A focus on conflation and the bidding costs at its root helps illuminate previously unexplained characteristics of these auctions.

Every Internet search triggers an automated auction, which determines the placement of advertisements on a search results page. The earliest search auction, initiated by Overture (then GoTo.com) in 1998, required that participants bid a separate price to place an ad in each position on the search page. A sequence of first-price auctions determined the placements, starting with an auction for the first position and proceeding downwards. A winning bidder for one position was eliminated from the bidding for lower positions on that page.

This initial auction design involved no conflation: advertisers offered a separate price for each search position. GoTo soon abandoned this design in favor of one with

¹⁰ Nearby frequency bands can vary substantially in value due largely to differences in interference problems.

conflation and with a different pricing rule: advertisers' bids and payments were expressed on a price-per-click basis. The highest bidder was placed in the first position, the second-highest in the second position, and so on, and each bidder paid its price only when its ad was clicked. Since ads placed in lower positions on the page attract fewer clicks, the per-click pricing provides some automatic quality-adjustment in prices paid.

In more recent search auctions,¹¹ the rule used to assign advertisers to position on the page has again changed several times. In the incarnation we will study, the per-click bid prices were converted to adjusted bids by multiplying them by an estimate of the *click-through rate* (CTR), which is the fraction of ad exposures on which the viewer clicks. The adjusted bid can be described as the *expected payment per impression*. Comparing adjusted bids allows the search engine to identify which bidder has offered to pay the most for real estate on the search page. The highest adjusted bid wins the top position, the second-highest wins the next position, and so on. In the so-called *generalized second-price auction*, the amount an advertiser pays when its ad is actually clicked is the smallest price per click that, if bid, would have won that ad position. Simple algebra shows that an advertiser's expected payment per impression is then equal to the next-highest adjusted bid, so it is sometimes convenient to analyze the auction as if bids were submitted directly in this adjusted form.

If the bidders and the seller know the click-through rates, the generalized second-price auction is strategically equivalent to a series of second-price auctions in which all bids are adjusted (per-impression) bids and separate prices are offered for each position,

¹¹ Google adopted the design described in the text in 2002. Yahoo! switched to adopt essentially the same design in 2007. Rules continue to evolve, so this description may not be current.

but in which the auction is conducted with two restrictions.¹² The first resembles the one used in earlier search-advertising auctions: an advertiser who wins one position on a page is excluded from bidding for the lower positions. The second is a conflation: an advertiser's per-click bid must be the same for every position, so its corresponding adjusted bid for the k^{th} position is scaled in proportion to the lower number of expected clicks for the k^{th} position.

Thus, generalized second-price auctions differ from Vickrey auctions in two ways. Generalized second-price auctions conflate: a bid in these auctions is a single number, while the bids in a standard Vickrey mechanism is a vector with dimension equal to the number of positions on the page. Also, given the vector of values implied by the one-dimensional bids and estimated click-through rates, the prices of the ad positions is determined differently from those of a Vickrey auction: they coincide instead with the prices in a certain series of second-price auctions. These two differences suggest corresponding questions: First, what effects does conflation have on the performance of the generalized second-price auction? Second, if the auctioneer must conflate because it is restricted a single, one-dimensional bid, what advantage might it gain by using a generalized second-price auction instead of using the conflated bid and click-through rates to infer a multi-dimensional bid and then applying the Vickrey pricing rule?

We analyze the effects of conflation using a model in which longer messages incur higher costs. For our model of sponsored search auctions, we assume that each *price* specified by a bidder incurs a positive cost. Such costs may be associated with

¹² In early postings describing the auction, Google claimed that this generalized second-price auction was the actual Vickrey auction, but that is a mistake. In particular, no bidder has a dominant strategy in the generalized second-price auction.

deciding how much to bid, submitting the bids themselves, and monitoring the bids to evaluate how they are performing and to check whether changes are needed. With these costs, conflation can sometimes be justified by the cost savings it creates, but its bigger advantage may be the effect it has on auction outcomes.

What is unexpected is that introducing arbitrarily small positive bid costs can lead to a profound change in the equilibrium outcomes of the standard (unconflated) Vickrey and generalized second-price auctions. The reason is that in both mechanisms, it is the losing bids for the various positions that determine the prices. At any pure-strategy equilibrium, if bid costs are positive, bidders make no losing bids, so there is no pure-strategy, full-information equilibrium, of either the Vickrey auction or the sequence of second-price auctions, at which the seller earns positive revenue.¹³

With positive bid costs, we return to the first question and find a valuable role for conflation: it eliminates the zero-revenue equilibria. Every equilibrium of the generalized second-price auction for two or more items generates positive revenues for the seller, because the bid by the winner of any position $n > 1$ establishes a positive price for the winner of position $n-1$.

The second question concerns the advantage of the generalized second-price rule compared with the Vickrey rule. Although the literature already includes analyses highlighting important disadvantages of the Vickrey pricing formula in multi-item auctions (Ausubel and Milgrom (2005), Rothkopf (2007)), the most devastating objections apply only when bidders can buy multiple items. In sponsored-search auctions,

¹³ This paper uses full-information Nash equilibrium to analyze various mechanisms. Based on earlier empirical successes and failures of game-theoretic auction models, what we believe should be taken most seriously from this analysis is the *comparative* predictions about the revenue performance of alternative auctions mechanisms, rather than the point predictions about the performance of any single mechanism.

each bidder acquires at most one position, so the main objections of the earlier literature have no force.

Our answer to this second question focuses on the special environment of sponsored search, extending the models used in earlier studies to allow a simple form of heterogeneity among searchers. For simplicity, we assume that there are two kinds of searchers—some are *shoppers* who are actually looking for a product to buy and others are *merely curious* about the products being advertised—with each group having its own click rates for ads occupying different positions on the page.¹⁴ Each advertiser has some positive value per click from shoppers, but a zero value per click from the merely curious. Also, the frequency of clicks from each group falls as one moves down the search page.

The differences in click rates between the shoppers and the merely curious are important to advertisers in formulating their bid strategies. For example, if shoppers look more carefully than others and click more frequently on ads near the bottom of a search page, then the value per click of those ads will be higher than for ads near the top of the page, because the lower ranking ads will include a higher fraction of shoppers. In general, if the click rates differ among groups with different propensities to buy, then clicks from different positions have different values.

In the model, bidders' value types are one-dimensional, so the value per impression declines as one moves down the page, just as in the prior literature. What distinguishes our analysis is that the bidders have other relevant information that the

¹⁴ A simple model with two types is sufficient to make the two central points of this section. First, heterogeneity destroys the ability of the auctioneer to implement a Vickrey auction from conflated bids. Second, a generalized second-price auction can sometimes achieve the Vickrey outcome with conflated bids when the Vickrey auction itself cannot.

auctioneer does not. Based on its own data, the auctioneer can infer the average click rate for each position, but not the purchase behavior of clickers once they leave the search page, and thus not the proportion of types, at each position. Consequently, the auctioneer cannot infer from a bidder's reported value for an ad in one position what the bidder's values are for ads in the other positions. Therefore, if bidders made one-dimensional reports, the auctioneer will have too little information to implement the Vickrey pricing rule accurately, so prices will vary from the Vickrey prices. In contrast, for the generalized second-price auction with heterogeneity among searchers, there is sometimes a full-information equilibrium in which the realized prices are Vickrey prices. This is possible because each bidder can observe what proportion of its own clicks from various ad positions convert into sales and profits. To summarize: *a generalized second-price auction using conflated messages may, in equilibrium, lead to Vickrey prices in a wider set of circumstances than a similarly conflated Vickrey auction.*

The rest of this paper is organized as follows. Section II states and proves the *Simplification Theorem*, which shows that, for general games, simplifications that restrict the strategy set while satisfying an *outcome-closure property* never result in additional pure-strategy Nash or ϵ -Nash-equilibrium profiles. Section III treats the generalized second-price auction of sponsored search. Its first subsection shows that the generalized second-price auction is a tight simplification of a series of second-price auctions and eliminates certain problematic zero-revenue Nash equilibria. Its second subsection extends the results of prior research by introducing a model with two types of searchers and demonstrating that the selected equilibrium of the generalized second-price auction can still establish Vickrey prices. Section IV concludes.

II. The Simplification Theorem

Let (N, X, ω) be a normal-form mechanism where N is the set of players; $X = (X_1, \dots, X_N)$ is the set of strategy profiles; $\Omega \subseteq \Omega_1 \times \dots \times \Omega_N$ is the set of possible outcomes, where each Ω_n is endowed with a topology; and $\omega: X \rightarrow \Omega$ is the outcome function. Given utility functions $u_n: \Omega_n \rightarrow \mathfrak{R}$, one obtains a normal-form game in which the payoff to player n is $u_n(\omega_n(x))$.

We assume throughout that the utility functions are continuous. The representation of the outcomes in product form provides a way for the analyst to formulate restrictions on the participants' preferences. For example, one can formulate the assumption that participant n in an auction cares only about its own allocation and price by specifying that ω_n consists of just n 's allocation and price. Alternatively, to formulate the less restrictive assumption that parties may care about the full allocation – their goods, others' goods, and all the payments – then one would set $\omega_1 = \dots = \omega_N$.

Definitions.

1. If $\hat{X} \subseteq X$ and $\hat{\omega} = \omega|_{\hat{X}}$, then the mechanism $(N, \hat{X}, \hat{\omega})$ is a *simplification* of (N, X, ω) , and (N, X, ω) is an *extension* of $(N, \hat{X}, \hat{\omega})$.
2. The simplification $(N, \hat{X}, \hat{\omega})$ of (N, X, ω) has the *outcome closure property* if, for every player n , every profile $\hat{x}_{-n} \in \hat{X}_{-n}$, every $x_n \in X_n$, and every open neighborhood O of $\omega_n(\hat{x}_{-n}, x_n)$, there exists $\hat{x}_n \in \hat{X}_n$ such that $\omega_n(\hat{x}) \in O$.

3. For $\varepsilon \geq 0$, a strategy profile $x \in X$ is an ε -Nash equilibrium of (N, X, ω) for utilities u if for each player n and every strategy $x'_n \in X_n$,
- $$u_n(\omega_n(x'_n, x_{-n})) \leq u_n(\omega_n(x)) + \varepsilon.$$
- A Nash equilibrium is an ε -Nash equilibrium for $\varepsilon = 0$.
4. The simplification $(N, \hat{X}, \hat{\omega})$ of (N, X, ω) is *tight* if, for every profile $u = (u_1, \dots, u_N)$ of continuous utility functions and every $\varepsilon \geq 0$, every pure strategy profile x that is an ε -Nash equilibrium of $(N, \hat{X}, \hat{\omega})$ is also an ε -Nash equilibrium of (N, X, ω) .

Theorem 1 (Simplification Theorem). Any simplification $(N, \hat{X}, \omega_{|\hat{x}})$ of (N, X, ω) that has the outcome-closure property is tight.

Proof. Fix $u \in U$ and $\varepsilon \geq 0$. Assume that the simplification $(N, \hat{X}, \omega_{|\hat{x}})$ of (N, X, ω) has the outcome-closure property and suppose that the profile \hat{x} is not an ε -Nash equilibrium of (N, X, ω) . Then, for some player n and strategy $x'_n \in X_n$,

$$\delta \equiv u_n(\omega_n(x'_n, \hat{x}_{-n})) - u_n(\omega_n(\hat{x})) - \varepsilon > 0.$$

Since u_n is continuous, by the outcome closure property there exists some $x''_n \in \hat{X}_n$ such that $u_n(\omega_n(x'_n, \hat{x}_{-n})) - u_n(\omega_n(x''_n, \hat{x}_{-n})) < \delta$, and hence $u_n(\omega_n(x''_n, \hat{x}_{-n})) - u_n(\omega_n(\hat{x})) - \varepsilon > 0$. Therefore, \hat{x} is not an ε -Nash equilibrium of $(N, \hat{X}, \omega_{|\hat{x}})$. ♦

The Simplification Theorem describes a property of the set of pure equilibria of a simplified direct mechanism. Since mixed equilibria are just pure equilibria of an

extended game, the theorem can be applied to study mixed equilibria as well. In general, outcome closure for mixed strategy profiles is more demanding than for pure strategy profiles.

III. Application to Search Auctions

In a pair of recent papers, Edelman, Ostrovsky, and Schwartz (2007) and Varian (2006) have studied the generalized second-price auction using the assumptions that bidders value all clicks equally (regardless of the position of the ad) and that a bidder's payoff is equal to the value of its clicks minus the total amount it pays. A central finding of both papers is that the prices and position assignments from a particular, selected, full-information Nash equilibrium of the generalized second-price auction are the same as those from the dominant strategy equilibrium of a Vickrey auction. The prices also coincide with the minimum competitive market-clearing prices, given the bidders' reported preferences.

This theory reveals interesting connections, but fails to explain much about why sponsored-search auctions are run as they are. First, why do advertisers pay on a per-click basis, rather than on the per-impression basis that is most commonly used for print, radio and television advertising? Second, given that only per-click bids are accepted, why determine the pricing of ad positions by a sequence of second-price auctions rather than the Vickrey formula?

The first question is puzzling only in a static, full-information environment, where per-click and per-impression pricing are just two ways of expressing the same total payment. In such a hypothetical setting, there would be little reason to prefer one over another. In practice, per-click charges are easier for an Internet advertiser to audit

because it can meter visits to its own website. Also, as discussed in the introduction, per-click pricing facilitates conflation in search bidding. Beyond that, search companies have continually expanded the scope of bids, showing ads on a wider variety of sites and encouraging advertisers to use “extended-match” technologies to place ads not only on pages that match the bidder’s search terms exactly, but also on pages that match them only approximately. For example, the extended-match technology might deem the term “ink cartridge” closely enough related to the term “printer cartridge” to show an ad for the latter when a search is made for the former. The relation among these search terms is, however, imperfect – for example, “ink cartridge” might be entered by a user searching for a pen-ink refill – so the proportion of searchers who are shoppers for a printer cartridge company may be lower for the related terms than for the exact ones, which makes each impression less valuable. Even click values may be different, because clicks from pen-ink searchers would less frequently result in actual sales. With such extensive conflation, pricing ads on a per-click basis is helpful because it reduces the advertiser’s cost per impression for less-valuable ads and makes advertisers more willing to use the extended match technology. The auctioneer has a limited supply of real estate on search pages, and basing prices on clicks thus shares attributes with like adjustments made in the auction prices of rough diamonds based on average quality in a slice, as well as with quality-based adjustments in the last stage of the UK 2.6 GHz auction. Both of these adjustments were discussed in the introduction.¹⁵

¹⁵ Bidders, especially so-called “performance advertisers,” who hope to convert clicks directly into sales, may have a different perspective, seeing themselves as buying customer visits, rather than impressions. This difference in perspectives does not change the analysis in any fundamental way.

How Bid Costs Affect Mechanism Performance

We turn now to the very damaging effects of positive bid costs on revenues in the Vickrey and sequential-second-price auction mechanisms and to an analysis of how conflation mitigates those effects.¹⁶

Suppose that bidder n 's value for an ad in position k is denoted v_{nk} . Each advertiser is permitted to acquire only one ad position, so the vector v_n completely describes the bidder's values for the possible positions it might acquire. We make the standard normalization that a bidder who gets no ad has a zero payoff. For this part of the analysis, it does not matter whether bids are made on a per-impression or a per-click basis, so we take them to be per-impression bids. Suppose that there is a small cost $\varepsilon > 0$ of submitting a positive bid for each placement or position on the search page.¹⁷ In this model, for any profile of bids made by competitors, a bidder's best reply never includes any positive losing bids, and the usual dominant strategy analysis for the Vickrey auction fails. One part of the dominant-strategy analysis does, however, have a useful counterpart in the model with costly bidding: if a bidder n submits a positive bid $v_{nk} \neq b_{nk} > 0$ for just one position, then that bid is weakly dominated by the bid $b_{nk} = v_{nk}$. By inspection, the strategy profile in which bidders bid $b_{nk} = v_{nk}$ for items efficiently assigned to them, and bid zero for other items, constitutes a Nash equilibrium. Notice that the same profile is a Nash equilibrium of the Vickrey auction even when bid costs are zero. Let us summarize:

¹⁶ Throughout our analysis of auctions, we ignore the possibility of ties. These can be treated by an extension of the equilibrium concept, as suggested by Simon and Zame (1990), or by other devices, but such details do not affect any substantive conclusions.

¹⁷ Bidders in sponsored-search auctions typically leave their bids in place across a large number of auctions, reviewing them only periodically. One can think of ε as representing the average cost-per-bid-submitted of the initial bid submission and the periodic reviews.

Theorem 2. In any pure-strategy equilibrium of the Vickrey auction with costly bidding ($\varepsilon > 0$), the seller's revenue is zero. If all bidders play weakly undominated strategies, then in equilibrium the winner n of position k bids $b_{nk} = v_{nk}$ for that position and zero for all other positions. This zero-revenue equilibrium bid profile is also a pure Nash equilibrium of the game when the bid cost is zero ($\varepsilon = 0$).

To make an analogous statement for a sequence of second-price auctions, we let the vector $b_n = (b_{n1}, \dots, b_{nK})$ denote the bids that advertiser n is prepared to make for each of the K positions. To keep notation simple, let us permute the bidder indexes so that bidder 1 is the bidder who wins the first position, bidder 2 wins the second, and so on. Let $L_n = \max_{j > n} b_{jn}$ denote the second-highest bid among the bidders for position n . In the sequence of second-price auctions, this is the price paid by bidder n to acquire ad position n . If bidder n makes J_n positive bids, then its payoff is $v_{nn} - L_n - \varepsilon J_n$.

Theorem 3. In any pure equilibrium of the sequence of second-price auctions with costly bidding ($\varepsilon > 0$), the seller's revenue is zero. In any pure equilibrium in weakly undominated strategies, the winner n of position k bids $b_{nk} = v_{nk}$ for that position and zero for all other positions. There is a zero-revenue, undominated equilibrium in which the items are assigned efficiently. This same bid profile is also a pure Nash equilibrium when the bid cost is zero ($\varepsilon = 0$).

In both the Vickrey auction and the sequence of second-price auctions, a suitable conflation eliminates the zero-revenue equilibria. For the Vickrey auction, we temporarily follow the earlier papers in assuming that bidders' values per click do not

depend on the ad position, and that the click rate on an ad in position k is some fixed fraction α_k of the rate for position 1, where $1 = \alpha_1 > \dots > \alpha_K > 0$. Then, bidder n 's vector of values for the K positions is $v_{n1}(1, \alpha_2, \dots, \alpha_K)$, which can be completely described by reporting the single number v_{n1} , henceforth denoted more simply by v_n . Consider a game in which, if bidder n bids b_{n1} for the first position, then bids of $b_n = b_{n1}(1, \alpha_2, \dots, \alpha_K)$ are imputed for the K positions. The auctioneer can use those bids to compute the imputed Vickrey prices for each bidder and position.

In the resulting game with positive bid costs, any bid $v_n \neq b_{n1} > 0$ is weakly dominated by the bid $b_{n1} = v_n$. If the bid cost ε is small but positive, then, in an undominated, pure equilibrium, each of the bidders with the K highest values will make positive bids and the other bidders will bid zero. The resulting prices will be the Vickrey prices that would emerge if just the K highest bidders had participated, and these prices are given by the usual formula (Edelman, Ostrovsky, and Schwartz (2007), Varian (2006)).

Theorem 4. Suppose there are $K > 1$ positions for sale, at least K bidders, and small positive per-bid costs $\varepsilon > 0$. Then, for every pure-strategy Nash equilibrium of the simplified Vickrey auction, the price paid for position K will be zero. If the bidders are ordered such that $v_1 > v_2 > \dots > v_K > \dots > v_N$, then the price of each position $k < K$ is

$$p_k = \sum_{i=k+1}^K (\alpha_{i-1} - \alpha_i)v_i > 0. \text{ Equilibrium revenues are strictly positive.}^{18}$$

A similar analysis applies to the generalized second-price auction.

¹⁸ For Theorems 4 and 5, if there are $N < K$ bidders, then the outcome is the same as if there were just N items for sale.

Theorem 5. Suppose there are $K > 1$ positions for sale, at least K bidders, and small positive per-bid costs $\varepsilon > 0$. Then, in every pure-strategy equilibrium of the generalized second-price auction, the price paid for position K will be zero, but all other prices will be strictly positive. Equilibrium revenues are strictly positive.

Even with zero per-bid costs ($\varepsilon = 0$), there are still zero-revenue Nash equilibria of the unconfated versions of the Vickrey auction mechanism and the sequence of second price auctions. These are the same as the equilibria for small positive ε . The generalized second-price auction, with its conflation of bids, eliminates the zero-revenue problem even when $\varepsilon = 0$, leaving intact only the equilibria with positive revenues.

The generalized second-price auction has the additional advantage that it satisfies the outcome-closure property and so, by the Simplification Theorem, avoids introducing any new equilibria or ε -equilibria. Suppose that, in any series of second-price auctions in which all advertisers other than j use one-dimensional strategies, bidder n alone is permitted to submit separate bids for each position. Any bid by n in this extended message space that wins position k pays a price equal to the $k+1^{\text{th}}$ highest bid, and n could obtain the same position at the same price by making the k^{th} highest one-dimensional bid. In view of the Simplification Theorem, we have proved the following:

Theorem 6. The generalized second-price auction has the outcome closure property relative to a sequence of second-price auctions in which bidders are restricted to winning only one position.

How Conflation Affects Comparative Performance

How does the use of conflation affect the comparative performance of alternative pricing rules? The previous literature has shown that, in one particular model, a selected equilibrium outcome of the generalized second-price auction and the dominant strategy outcome of the Vickrey auction are one and the same. If this outcome is desirable, that would seem to suggest the search companies ought to prefer the Vickrey mechanism to the generalized second-price auction, because it achieves the Vickrey outcome more reliably as a dominant-strategy solution.

It is obvious that if bidders differ only in their values per click and if there is no other private information, then the Vickrey mechanism can be implemented with one-dimensional reports, because types are then one-dimensional. To implement the Vickrey mechanism, one converts the type to a vector of position values by multiplying the reported value per click by the click-through rate for each position. Applying the Vickrey formula to the resulting vector of position values sets prices. If the generalized second-price auction does no more than achieve the same outcome using one-dimensional reports with a selected full-information Nash equilibrium, then why not use the Vickrey auction instead? Does the generalized second-price auction have some actual advantage compared to the conflated version of the Vickrey mechanism?

To explore these questions, we extend the preceding models to allow for searcher heterogeneity. To illustrate the effect of heterogeneity, it is sufficient to suppose that there are two types of searchers. Searchers of one type (“shoppers”) are looking for a product to buy while those of the other (“merely curious” searchers) are just looking for

information. The probability that a searcher is a shopper is $\lambda \in (0,1)$ and the probability the searcher is merely curious is $1 - \lambda$.

It is supposed in the literature that a searcher's click rate on an ad is determined by multiplying the ad's "clickability," which is the ad's click-through rate if it is placed in the first position on the search page, by the click-through rate for the position. Here, we assume the same, and also that the click-through rate for shoppers for an ad in position k is α_k ; for the merely curious, it is β_k . We assume that $\alpha_1 > \dots > \alpha_K > 0$ and $\beta_1 > \dots > \beta_K > 0$, but we do not assume that the two series are proportional. For example, if the attention of merely curious flags more quickly than that of shoppers, then the sequence β_k / α_k would be decreasing.

We assume that only clicks by shoppers are valuable to advertisers, so the value of an ad in position k to bidder n is $\lambda \alpha_k v_n$. A bidder can learn this positional value over time by observing its sales from ads in position k . This set of assumptions implies that assortative matching is efficient, that is, that the advertiser with the highest value v_n should be shown in first position, and so on for the other positions. Let us label the bidders so that $v_1 > \dots > v_N$ and assume that there are weakly more bidders than positions, $N \geq K$. Then, at the efficient allocation, position n is assigned to bidder n .

The click-through rate for position k is $\gamma_k = \lambda \alpha_k + (1 - \lambda) \beta_k$. Although this rate decreases with k , it would be a rare coincidence for it to decrease in direct proportion to the rate of valuable clicks, $\lambda \alpha_k$. If the search company observes clicks but not sales or value for each position, its auction rule can entail adjusting bids in proportion to clicks

but not in proportion to value. So, if bidder n names a price of b_{n1} for position 1 in a simplified auction, then the auctioneer would impute a bid of $b_{n1}\gamma_k / \gamma_1$ for position k , but that is not the same as the bidder's private value and *cannot* be used to infer accurate Vickrey prices.

We asked the question: why not use a one-dimensional Vickrey auction? The answer is that when there is bidder heterogeneity, it is not possible for the auctioneer to run an accurate Vickrey auction using one-dimensional bids. But the answer cannot explain the continued use of the generalized second-price auction unless that mechanism has good properties for at least some environments with bidder heterogeneity. We investigate that next.

It has been known for more than half a century that market-clearing prices exist for a class of matching problems which includes the heterogeneous bidder model; that there is a unique minimal market clearing price vector p for the positions; and that the equilibrium can be computed using linear programming (Koopmans and Beckmann (1957)). The minimum equilibrium price p_n is the shadow price of an additional impression in position n . It follows that p_n is the opportunity cost of the ad placed in position n by bidder n , so it is also the Vickrey price paid by bidder n to acquire that position.

Competitive equilibrium prices satisfy the constraint that bidder n prefers position n to position $n-1$, that is, $\lambda v_n \alpha_n - p_n \geq \lambda v_n \alpha_{n-1} - p_{n-1}$. And the single-crossing structure of preferences built into this expression ensures that these hold as equalities at the minimum competitive equilibrium: $\lambda v_n \alpha_n - p_n = \lambda v_n \alpha_{n-1} - p_{n-1}$. Defining $\alpha_{K+1} = 0 = p_{K+1}$, it

follows that the Vickrey/competitive equilibrium prices are

$$p_k = p_{K+1} + (p_K - p_{K+1}) + \dots + (p_k - p_{k+1}) = \lambda \sum_{i=k}^K (v_{i+1}(\alpha_i - \alpha_{i+1})).$$

This is also the formula for Vickrey prices reported by Edelman, Ostrovsky, and Schwartz (2007).

Theorem 7. For the model with two types of searchers, the following two conditions are equivalent:

1. There is a pure-strategy Nash equilibrium of the generalized second-price auction in which the assignment is efficient and prices paid by the winning bidders are the Vickrey prices p .
2. The Vickrey price-per-click sequence $\{p_k / \gamma_k\}_{k=1}^K$ is decreasing.

Proof. Order bidders so that $v_1 > v_2 > \dots > v_N$. An efficient allocation requires that bidder n win position n . To get the Vickrey prices for all positions, it is necessary that the equilibrium bid by bidder $n \in \{2, \dots, K+1\}$ be $b_{n1} = p_{n-1} / \gamma_{n-1}$. Take $b_{11} = \alpha_1 v_1$ and $b_{k1} = 0$ for $k > K+1$.

First we prove that if 2 does not hold, then 1 does not hold. If the Vickrey price-per-click sequence p_n / γ_n is not decreasing, then the bids are not ranked in the order required for an efficient assignment. (For example, if $p_3 / \gamma_3 < p_4 / \gamma_4$, then bidder 4 must be bidding less than bidder 5. The resulting assignment is inefficient because bidder 5, who has the lower value, receives a higher position than bidder 4.)

Next suppose that $\{p_n / \gamma_n\}_{n=1}^N$ is decreasing, and fix any bidder n . Recall that the Vickrey prices are competitive equilibrium prices, so no bidder wishes to deviate in order to purchase a different position at prices p . If bidder n raises its bid to win a higher

position, say position $k < n$, then the price it must pay is determined by the k^{th} highest opposing bid, so it is $\gamma_k(p_{k-1}/\gamma_{k-1}) > \gamma_k(p_k/\gamma_k) = p_k$, making that deviation unprofitable. If bidder n reduces its bid to win a lower position $k > n$, then the price it must pay is precisely p_k , and the deviation is again unprofitable. ♦

IV. Conclusion

In the theory of mechanism design as in the design of physical mechanisms, intuition tells us that simplicity can be beneficial. But the intuition is vague. Defining simplicity and identifying its advantages pose challenges. In this paper, we have defined a *simplified mechanism* to be a mechanism in which the set of messages or reports has been restricted. Well-conceived simplifications of direct mechanisms can reduce costs and eliminate certain undesirable equilibria.

We have explored conflation, a common simplification for auction mechanisms, according to which two or more distinguishable outcomes are conflated by the auction rules so that the same bids apply to both. Sometimes, conflation is so obviously advantageous that its presence goes entirely unnoticed. For example, in Treasury auctions, bills with different serial numbers are conflated: it would be *possible* to sell such bills separately. Are there advantages of doing so? On the contrary, in the Treasury bill example, depending on the specification of bidding costs, simplification can eliminate not just undesirable pure equilibria, but undesirable mixed equilibria as well.¹⁹

¹⁹ Suppose there are N bills and $N+1$ bidders, and that each bill is worth 1 to each bidder. Introduce bid costs by assuming that each bidder can bid costlessly for one bill, but incurs a cost to bid for two or more. Then, the unique Nash equilibrium of the simplified first-price mechanism with a zero minimum bid has revenue of N . However, no pure or mixed equilibrium of the unconflicted auction for N individual items has expected revenue greater than 1.

The next step in this research program is to apply the principles of simplification to invent and evaluate new mechanisms. The first such project is a sequel to this paper, Milgrom (2008), which introduces a new message space for expressing substitutable preferences, and a corresponding simplification of the Walrasian exchange mechanism that satisfies the outcome closure property.

References

- Ausubel, Lawrence and Paul Milgrom. 2005. "The Lovely but Lonely Vickrey Auction," in *Combinatorial Auctions*. Peter Cramton, Yoav Shoham and Richard Steinberg eds. Cambridge, MA: MIT Press.
- Bernheim, B. Douglas and Michael Whinston. 1986. "Menu Auctions, Resource Allocation and Economic Influence." *Quarterly Journal of Economics*, 101, pp. 1-31.
- Clarke, E.H. 1971. "Multipart Pricing of Public Goods." *Public Choice*, XI, pp. 17-33.
- Edelman, Benjamin, Michael Ostrovsky, and Michael Schwartz. 2007. "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords." *American Economic Review*, 97:1.
- Gale, David and Lloyd Shapley. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly*, 69, pp. 9-15.
- Groves, Theodore. 1973. "Incentives in Teams." *Econometrica*, 61, pp. 617-31.
- Hatfield, John and Paul Milgrom. 2005. "Matching with Contracts." *American Economic Review*, 95:4, pp. 913-35.
- Kelso, Alexander and Vincent Crawford. 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica*, 50:6, pp. 1483-504.
- Koopmans, Tjalling and Martin Beckmann. 1957. "Assignment Problems and the Location of Economic Activities." *Econometrica*, 25:1, pp. 53-76.
- Milgrom, Paul. 2008. "Assignment Messages and Exchanges." *Working paper*.
- Roberts, John and Andrew Postlewaite. 1976. "The Incentives for Price-Taking Behavior in Large Exchange Economies." *Econometrica*, 44:1, pp. 115-29.

- Roth, Alvin E. and Elliott Peranson. 1999. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review*, 89, pp. 748-80.
- Roth, Alvin E. and Marilda Sotomayor. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge: Cambridge University Press.
- Rothkopf, Michael. 2007. "Thirteen Reasons Why the Vickrey-Clarke-Groves Mechanism is Not Practical." *Operations Research*, 55:2, pp. 191-97.
- Simon, Leo K. and William R. Zame. 1990. "Discontinuous Games and Endogenous Sharing Rules." *Econometrica*, 58, pp. 861-72.
- Varian, Hal R. 2006. "Position Auctions." Working paper.
- Vickrey, William. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, XVI, pp. 8-37.