

Should the Treasury Price-Discriminate? A Procedure for Computing Hypothetical Bid Functions

by

DANIEL HELLER AND YVAN LENGWILER*

Ever since FRIEDMAN [1959], it has been debated whether the Treasury should sell bonds in a uniform or in a discriminatory auction. Empirical research on this topic has been confined to experiments in which both auctions were used. But these experiments inherently contain an identification problem, since differences cannot be attributed to the auction format alone. We develop a method for generating counterfactual data on discriminatory auctions, using real data from uniform-price Treasury bond auctions in Switzerland. Our method allows us to investigate the performance of the two auctions without relying on experiments and without the identification problem. (JEL: D 44, H 63)

1 Introduction

For decades, governments have been auctioning off fixed-income securities to finance their debt. Despite the importance of Treasury auctions in terms of value, so-called multi-unit auctions have so far received fairly little academic attention. While auction theory has been a prolific field in economics, it has focused mostly on single-unit auctions such as auctions for a painting, a contract, or drilling rights to an oil field. Multi-unit auctions are different in that the good being auctioned is divisible. Starting with WILSON [1979], it has been shown in a number of papers that the conclusions reached in single-unit auction theory do not necessarily carry over to multi-unit auctions, see BACK AND ZENDER [1993], NOUSSAIR [1995], TENORIO [1997], LENGWILER [1998], ENGELBRECHT-WIGGANS AND KAHN [1998a], [1998b].

* We wish to thank V.V. Chari, Elisabeth Coutts, Ed Green, Michael Gordy, Peter Kugler, Dieter Nautz, Georg Nöldeke, Andy Sturm, Walter Wasserfallen, Elmar Wolfstetter, and the referees, as well as seminar participants at the Virginia Polytechnic Institute, at the Federal Reserve Bank of Minneapolis, and at the University of Basel for interesting comments. Parts of this paper were written while Yvan Lengwiler was on leave at the Board of Governors of the Federal Reserve System. The views expressed are not necessarily those of the Board of Governors of the Federal Reserve System or of the Swiss National Bank.

In practice, Treasuries apply two kinds of formats for auctions of fixed-income securities, namely discriminatory and uniform-price auctions. In a discriminatory auction a bidder pays the bid price if the bid price is above or equal to the cutoff price set by the Treasury. In a uniform-price auction, a bidder pays the cutoff price conditional on having submitted a bid that was equal to or higher than the cutoff price. In most countries, bidders are allowed to submit multiple price–quantity pairs.

Considerable attention has been given to the question of which auction format yields greater revenues for the Treasury. BIKCHANDARI AND HUANG [1993] provide a survey of this controversy. Until recently, discriminatory auctions for bonds were used more often than uniform auctions, but many countries have been experimenting with their auction formats. Until now, quantitative research on the question of which format is more lucrative for the Treasury has used data from periods in which the Treasury applied the two different auction formats, either simultaneously or sequentially. Yet, all of these experiments suffer from a fundamental identification problem, since observed differences may not be wholly attributable to the auction format.¹

In this paper, we present a computational procedure for transforming the aggregate bid function of a uniform-price auction into a corresponding bid function that would have been submitted under price discrimination. This allows us to compare the two auction formats even if no (quasi- or even true) experiments are available. Thus, we avoid the identification problem inherent in the experiments done so far.

As an illustration of our method we apply the procedure to the Swiss data. This allows us to determine which auction rule would have been more advantageous for the Swiss Treasury. We identify a statistically significant advantage of the uniform over the discriminatory auction format.

2 Constructing Hypothetical Bid Functions

2.1 Theory

An early model that is especially relevant for Treasury auctions is SMITH [1966]. NAUTZ [1995a] further developed this model by deriving the equilibrium bid functions. Our method for computing hypothetical bid functions builds upon this work, so we will briefly present the Smith–Nautz model.

Let there be a finite *price grid* $P := \{p_1, \dots, p_j, \dots, p_k\}$. Prices are indexed in increasing order, $0 < p_1 < \dots < p_j < \dots < p_k$. Every bidder i submits a *bid function* $b_i: P \rightarrow \mathbb{R}_+$ indicating the amount of debt he wishes to purchase at each price

¹ Good experimental design requires random assignment to experimental conditions that are the same in every regard except for the manipulation of interest. To date, this requirement has been violated by all Treasuries that have carried out trials with auction formats. Hence, it would be more appropriate to describe these tests as quasi-experiments than as experiments. See section 2.3 below for a more extensive discussion of this identification problem.

on the grid P . From these bids, the *aggregate bid function* is defined as $B(p) = \sum_i b_i(p)$. It is important to note that the submitted bids may deviate from true willingness to pay, since bidders do not have to reveal their true preferences. Let $d_i: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ denote the true demand function.² We assume that d_i is a strictly decreasing, invertible function. As before, $D(p) = \sum_i d_i(p)$. After the auction, the seller announces a *cutoff price* π . All bids associated with prices weakly greater than π will be met, so $B(\pi)$ is the number of bonds sold to the bidders.

In a *uniform-price auction*, all accepted bids are fulfilled at the cutoff price. For each unit of debt bidders purchase, they pay only the cutoff price, even if their bid price is greater. Thus, the government's revenue is simply

$$(1) \quad R(\pi) = \pi B(\pi).$$

Note that this is just like the revenue function of an ordinary monopolist. The government chooses the cutoff price in such a way that it can raise the amount of funds that it wants to. Let S denote this amount. In general, there will be more than one cutoff price π that satisfies this condition. For instance, if B is not "too convex," then R is single-peaked, and there will be two π 's that yield the targeted amount of funds. In this case, the government chooses the higher price, because this is associated with the issue of fewer government bonds.³

In a *price-discriminating auction*, a bidder pays the price at which he submitted a bid, even if the cutoff price is lower. To distinguish between the equilibria of the two auction formats, a tilde will denote the endogenous variables in a price-discriminating auction only. Accordingly, the government's revenue is

$$(2) \quad \tilde{R}(\tilde{\pi}) = \sum_{p_j \geq \tilde{\pi}} p_j (\tilde{B}(p_j) - \tilde{B}(p_{j+1})),$$

with $\tilde{B}(p_{k+1}) = 0$.

When submitting their bids, the bidders do not know the cutoff price that will result. This uncertainty is captured by some distribution function F from which π is drawn (or in the case of a discriminatory auction, $\tilde{\pi}$ is drawn from \tilde{F}). It is assumed that bidders are individually negligible and therefore none of them, on its own, can affect the seller's chosen cutoff price. As a consequence, bidders take F or \tilde{F} as given. Let $f(p_j) = F(p_j) - F(p_{j-1})$ for $j > 1$ and $f(p_1) = F(p_1)$, and likewise for the distribution of the price in the discriminatory auction.

Within this setup, NAUTZ [1995a] characterizes optimal bidding in both auction formats. With a few slightly different assumptions, Nautz's result can be somewhat simplified. First, we assume that all prices on the price grid have some probability of being chosen as the cutoff price, i.e., $f(p_j) > 0$ and $\tilde{f}(p_j) > 0$ for all j .

² In other words, d_i is the inverse function of the marginal willingness to pay.

³ There is a maximum of funds that can be raised, given a bid function, namely $S := \max \pi B(\pi)$. If the target S exceeds S , then no cutoff price can satisfy the government's financial needs.

This is, in fact, not a strong assumption, because we can always fix the price grid in such a way that it is contained in the support of the distribution of the price. Second, we do not restrict bidders to using decreasing bid functions. This assumption allows us to use simple unconstrained maximization. At least for the empirical application we provide, this, too, is not a very strong assumption, because, as we will see in section 3.4, this constraint actually never binds empirically. These assumptions imply the following simplification of Nautz's result.

Proposition 1 (NAUTZ [1995a]): In a uniform-price auction, each bidder i submits his true demand function as a bid, $b_i(p_j) = d_i(p_j)$, for all j . In a price-discriminating auction, bidder i 's bid function is given by

$$(3) \quad \tilde{b}_i(p_j) := d_i \left(p_j + (p_j - p_{j-1}) \frac{\tilde{F}(p_{j-1})}{\tilde{f}(p_j)} \right) < d_i(p_j)$$

for all $j > 1$, and $\tilde{b}_i(p_1) = d_i(p_1)$.

The following intuition is behind this result. Consider the discriminating auction first. Price discrimination can be seen as the government's attempt to extract the consumer surplus. Yet, bidders will try to avoid the extraction of this surplus and will behave differently in a price-discriminating auction than they do in a uniform-price auction (SMITH [1966]). According to equation (3), the bid shading that can be observed in a discriminating auction is the result of precisely this attempt: Bidders try to save some of the consumer rent by bidding weakly at high prices.

Bidding is more straightforward in the uniform-price auction. Consider a bid that is different from the true demand at some price p_j . If π turns out to be different from p_j , the bid amount is irrelevant for this outcome. Yet, if π turns out to be equal to p_j , the bidder will not receive his utility-maximizing quantity, which is $d_i(p_j)$, but the quantity he bid, $b_i(p_j)$. So the bidder does not stand to gain by misrepresenting his demand. As a consequence, he will optimally reveal his true willingness to pay.

This result is in sharp contrast to the collusive equilibria under uniform pricing found by WILSON [1979] and further developed by BACK AND ZENDER [1993] and WANG AND ZENDER [1999]. In their models, bidders submit a steeper demand function than their true preference in order to "make room for weaker bidders" and thus yield a lower cutoff price. But this result hinges crucially on the assumption that the Treasury will auction off a quantity of bonds that is known in advance.

This assumption is, however, inappropriate for Swiss Treasury auctions, since the government has the right to adjust the supply of bonds after having examined the bids. This uncertainty about the supply of bonds in an auction has recently attracted a considerable research effort. Several studies find that the highly collusive equilibria are not robust to supply uncertainty. The ability of the seller to adjust the supply to the bidding mitigates collusion (NYBORG [1997], LENGWILER [1998], BACK AND ZENDER [1999], PAVAN, LI CALZI, AND GILLI [1999], MCADAMS

[2000]). Thus, it is doubtful that the collusive equilibria are relevant to our application. LENGWILER [1998] considers a model which is similar to the Smith–Nautz model, except that no perfect competition is assumed, i.e., bidders do not treat the distribution of the cutoff price as given. Moreover, it is assumed that the seller can adjust the supply elastically after knowing the bids. Interestingly, the equilibrium bid functions in this model approach the result presented above in proposition 1 as we increase the number of bidders indefinitely. Thus, we may interpret the Smith–Nautz model as the limit of a multi-unit auction with perfectly elastic supply (and therefore supply uncertainty) and “almost competitive” bidders.

2.2 The Treasury’s Decision Problem

Since the purpose of a bond auction is raising funds to finance debt, the Treasury should choose the auction format that enables it to do so at the lowest possible cost. The cost of paying back the debt is proportional to the number of bonds issued. Therefore, it is optimal for the Treasury to choose the auction format that provides the same revenue with the smaller number of bonds. In other words, the government’s choice is $\min \{B(\pi), \tilde{B}(\tilde{\pi})\}$, subject to $R(\pi) \geq S$ and $\tilde{R}(\tilde{\pi}) \geq S$, with S being the amount of funds raised in the auction. Which format is better depends, of course, on the bid functions B and \tilde{B} . More specifically, it depends on the way agents shade their bids in the discriminating auction. Suppose bidders reported their true willingness to pay in a price-discriminating auction, i.e., they did not shade (in this case, $B = \tilde{B}$). Then the price-discriminating auction would clearly be better for the government, because it would offer a means of collecting the entire consumer surplus. Yet, as proposition 1 establishes, bidders do shade, and in fact, they shade precisely in order to keep some of the consumer surplus. Hence, choosing between uniform pricing and price discrimination involves a trade-off for the government. In the uniform-price auction, bidding is strong, but the government cannot collect the consumer surplus. In a price-discriminating auction, the government collects at least some of the rent, but bidding is weak due to shading. From a theoretical point of view, it is not possible to determine *a priori* if the advantage of rent collection outweighs the loss due to shading. Hence, we will investigate the question empirically.

2.3 Empirical Comparison of the Two Formats

The best and probably most accurate way to compare auction formats empirically would be an auction in which the participants do not know which format the Treasury will choose. Thus, they would be required to submit two bid functions, one for each format. By comparing these bid functions, we would then be able to find the best format for the Treasury. To our knowledge, however, no such experiment has ever been performed.

The experiments that Treasuries have been performing fall into one of two categories. In one sort of experiment, the Treasury first performs a series of auctions

using one format, and then switches to the other. The *Bundesbank*, for instance, followed this route in 1988 when it switched from uniform-price to discriminatory auctions for its repo contracts (NAUTZ [1995b]), Zambia made the same move with foreign exchange auctions (TENORIO [1993]), and Mexico switched from price discrimination to the uniform-price format for the auctions of its government debt (UMLAUF [1993]). The other sort of experiment is to use the two auction formats simultaneously, but to use them for securities with different characteristics. This is the approach of the Norwegian Treasury (BERG et al. [2000]). In Norway, government bonds are sold in uniform-price auctions, but government bills are sold in price-discriminating auctions. The U.S. Treasury used to apply discriminatory auctions only, but it did a small experiment auctioning some securities with uniform-price auctions in the 1970s (SIMON [1994]), and it repeated this experiment on a much larger scale in the 1990s (MALVEY, ARCHIBALD, AND FLYNN [1995], NYBORG AND SUNDARESAN [1996], MALVEY AND ARCHIBALD [1998]).⁴

However, both types of experiments have the same problem. They do not reveal information on the outcome of a particular auction that would have occurred under a different format. Since in one case (the German, Mexican, and Zambian experiments) the bid functions are submitted at different points of time, we do not know whether the differences can be attributed to the auction format or to some unknown exogenous shock. The results are also not comparable if price discrimination is used for one kind of fixed-income securities and uniform pricing for another kind (as in the U.S. and Norwegian experiments). Empirically observed differences do not necessarily stem from the auction format; they may, for instance, also be caused by shocks to the preferences for different maturities.

Here, we propose an indirect computational method that allows us to empirically compare the two common multi-unit auction formats. This method is based on the theoretical considerations laid out in section 2.1; its core consists of the derivation of hypothetical bid functions that would have occurred if the other auction format had been used. Our data cover only uniform-price auctions, because this is the only format used by the Swiss Treasury. Our strategy is to use these data to generate the distribution \tilde{F} of the cutoff price $\tilde{\pi}$ if price discrimination had been used. Once we have \tilde{F} , we can apply proposition 1 and transform the observed bid function, B , into a hypothetical one, \tilde{B} , describing how bidders would have behaved if the Treasury had been price-discriminating. Finally, we compare the observed cost that the government actually incurred in financing its debt with the cost that it would have had to bear had it used the discriminatory auction.

⁴ Recently, the U.S. Treasury has decided to discontinue the price-discriminating auction format. All federal debt is now auctioned using the uniform-price format. Milton Friedman's suggestion has finally been adopted, more than 40 years after he made it!

2.4 Estimating \tilde{F} and Constructing \tilde{B}

In the theoretical sections above, we have seen that the density function \tilde{F} plays a crucial role in determining how agents shade their bids in a discriminatory auction. What do we know about \tilde{F} ? Because different Treasury bonds are close substitutes, the price of a bond in a primary auction must in some way be related to the prices of the bonds that are traded in the secondary market. It seems reasonable to assume that every bidder closely follows the prices on the secondary market to extract a signal about the cutoff price that is about to be set by the Treasury. A bidder is not going to submit a bid at a price that he considers too high relative to what he would have to pay for a similar bond in the secondary market. In this way, the bid functions in the auctions will tend to shift depending on secondary-market conditions. One way to capture this nexus is

$$(4) \quad \pi_z = \alpha + \beta \pi_z^s + \varepsilon_z,$$

with z indicating one of the auctions of the sample.⁵

The cutoff price π is equal to the sum of three components. The constant α is a price wedge between primary and secondary market. The price π^s is the price at which a bond with the same characteristics (same time to maturity, same coupon, same callability) is traded on the secondary market just before the auction takes place. The final price component is an error term ε ; it stems from the Treasury's discretion, which is unknown to the bidders. In addition to the secondary market, a bidder could also use other information to extract a signal about the cutoff price – for instance, the prices in the when-issued market. But all this information should already be contained in the secondary-market prices. Thus, equation (4) should give us an accurate picture of the expectations of the agents, and the error term ε in this equation should give us an estimate of the extent of uncertainty.

According to proposition 1, the bid functions observed in the Swiss uniform-price auction, B_z , are the true demand functions. If we knew the probability distribution of the cutoff price chosen by the Treasury under a price-discriminating regime, \tilde{F}_z we could apply (3) and compute the shaded bid functions \tilde{B}_z . However, since there have been no price-discriminating Treasury bond auctions in Switzerland, we cannot compute \tilde{F}_z directly.

We proceed as follows: We estimate (4) based on data from all auctions z . Assuming a normal distribution of the error term, we take the estimated standard deviation of ε , as well as the estimates for α and β to generate F_z (one for each auction z). Using (3), we then compute the bid function \tilde{B}_z^1 that would have been realized in the price-discriminating auction if the bidders expected $\tilde{\pi}_z$ to be distributed according to F_z . Next, we take \tilde{B}_z^1 to determine the cutoff price $\tilde{\pi}_z^1$, which is the cutoff price that would have raised the amount S_z of funds. Of course, $\tilde{\pi}_z^1$ is not

⁵ Because the variance of prices on financial markets is not constant over time, there is a potential heteroskedasticity problem. This problem is not easy to address here, because the time between observations is not constant.

distributed according to F_z , so the two are not compatible. Hence, we need to iterate this procedure in the following way:

$$\begin{aligned} \pi_z &= \alpha^0 + \beta^0 \pi_z^s + \varepsilon_z^0 \Rightarrow F_z \Rightarrow \tilde{B}_z^1 \Rightarrow \tilde{\pi}_z^1, \\ \tilde{\pi}_z^1 &= \alpha^1 + \beta^1 \pi_z^s + \varepsilon_z^1 \Rightarrow \tilde{F}_z^1 \Rightarrow \tilde{B}_z^2 \Rightarrow \tilde{\pi}_z^2, \\ &\vdots \\ \tilde{\pi}_z^n &= \alpha^n + \beta^n \pi_z^s + \varepsilon_z^n \Rightarrow \tilde{F}_z^n \Rightarrow \tilde{B}_z^{n+1} \Rightarrow \tilde{\pi}_z^{n+1}, \\ &\vdots \end{aligned}$$

As this process converges, it leads to pairs of bid functions and distributions, $(\tilde{B}_z^n, \tilde{F}_z^n)$, that are more and more compatible with each other. In the limit, it gives us the outcome of hypothetical price-discriminating auctions.⁶

We stop the iteration when

$$(5) \quad \max_z \left| \tilde{B}_z^{n+1}(\tilde{\pi}_z^{n+1}) - \tilde{B}_z^n(\tilde{\pi}_z^n) \right| < \delta.$$

The left hand side is the maximum absolute change in the number of bonds necessary to raise an amount S_z , across all auctions z . We keep on iterating as long as this maximum change exceeds some small δ .

3 Application: The Auction of Swiss Treasury Bonds

3.1 The Auction Procedure for Swiss Government Bonds

The Swiss Treasury has always only used the uniform-price format in its auctions. The following paragraphs describe the auction design that is applied by the Swiss Treasury for bonds in more detail.

The auctions for bonds take place on a bimonthly basis. Shortly before the auction, the Treasury announces the characteristics of the bonds that are to be issued, such as time to maturity, coupon, and callability. The Treasury also states the maximum number of bonds that will be issued. Usually, however, much fewer bonds than this maximum number are sold in the auction. The Treasury also reserves the right to cancel an auction if it does not consider the bids satisfactory.

Like other Treasuries, the Swiss Treasury frequently chooses not to issue a newly designed bond, but instead decides to extend the volume of a previously issued series (reopening). In such cases, the new bonds have exactly the same characteristics (coupon, maturity, callability) as a previously issued series. For reasons explained later, we will use only the data on such reopenings.

⁶ Note that the reverse procedure is much simpler. If data from a price-discriminating auction are available, (3) can be used directly to compute the true demand and therefore the equilibrium bid function of a uniform-price auction (subject to the validity of the Nautz model). No iteration is needed in this case. HORTACSU [2000] pursues this strategy in a Wilson-type model, which is a common-value model with fixed supply.

The bidders are invited to submit as many price–quantity bids as they wish. The price components of the bids have to be on a five-cent grid. In addition to submitting price–quantity pairs, bidders may place pure quantity bids without a price. These unpriced bids have to be less than 100,000 Swiss francs each. Unpriced bids are relatively unimportant. They account for about 7.2% of the total amount sold in an average Treasury bond auction.

The bids have to be in by 12 o'clock noon on the day of the auction. Until spring 2001, the bids were submitted by fax; then the Treasury switched to a proprietary electronic platform. After all the bids have been submitted, the Treasury decides on the cutoff price. All bids greater than this price, as well as all unpriced bids, are fully satisfied. Since the Treasury also sets an upper bound on the number of bonds it will issue, rationing of bids submitted at the lowest winning price may occur. In recent years, however, rationing has become less frequent, since the announced upper bounds have been substantially increased.

The circle of bidders is restricted to institutions holding accounts with the Swiss National Bank (SNB), currently some 400 entities. Under the current regulation, these are all banks based in Switzerland. All bidders are treated equally, i.e., there are no primary dealers. Also, no fee is charged to the participants.

3.2 *The Data Set*

The characteristics of Treasury bonds that have been issued in Switzerland in the last few years vary widely: Times to maturity range from 2 to 50 years, most bonds are not callable, some include one or several call options. Often, there is no bond with exactly the same characteristics traded on the secondary market at the time of the auction. One way to deal with this problem is to first estimate a zero-coupon yield curve and then to use this curve to determine the price that the to-be-auctioned bond would obtain on the secondary market. This approach would allow us to correct for different maturities and coupons, but it fails if some of the bonds are callable.

In order to avoid this problem, we consider only reopenings and exclude initial public offerings. A bond auctioned in a reopening has exactly the same coupon, the same date of maturity, and the same call options as a previously issued bond that is still traded on the secondary market.⁷ We operationalize π^s as the secondary-market price of such a previously issued bond at 11 o'clock a.m. on the day of the auction, just one hour before the bidding window closes. This price should contain almost all the public information available about the value of the to-be-issued bond. Our sample covers all reopenings between October 1993 and June 2000, a total of 56 observations.

⁷ Reopenings are especially popular in smaller countries, since they deepen the market and increase its liquidity.

3.3 Application of the Nautz Model to Swiss Treasury Auctions

Our method for computing the hypothetical bid functions under price discrimination is subject to the validity of NAUTZ's [1995a] model. This model builds on two key assumptions. One is that each bidder is individually small, and thus does not take the effect of his behavior on the seller's decision into account. The other is that the demand functions are independent of each other, indicating that the bidders' valuations for Treasury bonds are private values. We address these assumptions in the above order.

The assumption of competitive behavior may look like a tall order in the Swiss market: It is true that between 20 and 30 bidders usually participate in the auctions, but despite this significant number, a few large Swiss banks typically buy between two-thirds and three-quarters of each issue. Nonetheless, there are three good reasons to consider the Swiss Treasury auction competitive.

First, previously issued government bonds are actively traded on the secondary market, and they provide close substitutes for the bonds that are being issued with the auction. This is especially true for reopenings, because exactly the same security is already available on the secondary market.

Second, the market is highly contestable, as the only requirement for participating is to have the status of a bank. But many banks do not usually participate in the auctions. Thus, if established bidders were able to extract large rents from the auctions, it would be worthwhile and very easy for the banks that typically do not participate to outbid the established participants. Thus, the scope for oligopolistic behavior is limited. In fact, no episodes of market cornering or collusion are known in the Swiss Treasury market.

Third, as we mentioned above, the ability of the seller to adjust the supply in response to the bidding strongly works against collusion. The fact that the Treasury does not *ex ante* commit to a specific quantity severely constrains the bidders' ability to affect the Treasury's chosen cutoff price. If bidding is weak, the Treasury is not forced to accept a low cutoff price. Since the Swiss Treasury always tends to hold a substantial amount of spare liquidity, it can simply reduce the number of bonds issued in such a case. In fact, the Treasury even has the discretion to cancel an auction altogether if it deems the bids unsatisfactory – an event that has occurred once in the last six years.

The assumption of private values, which is built into the Nautz model, may also seem problematic at first sight. Private values are not an obvious choice for a model of a financial market. After all, many bidders buy in order to resell. Consider a slightly more elaborate description of the demand functions featured in the model used in section 2.1. We can write them as $d_i(p, m, s_i)$, where m is some public information about market conditions (and thus about the value of the government bond), and s_i is some private information. The function d_i captures the (true) demand of investor i as a function of these public and private signals, and of the price that has to be paid for the bonds. We argue that the private signals are of minor importance in our context, simply because the public information that is avail-

able is so good, especially for the reopenings that we use in the estimation. More specifically, the secondary market price of the to-be-issued bond can definitely be considered to be public information. In fact, let us assume that all the public information, m , is contained in the secondary-market price, so $m: = \pi^s$. This publicly observable signal contains a great deal of information about the resulting price of the auction, as can be established by regressing π^s on π (estimating (4)). This regression yields an adjusted R^2 of 99.4%.⁸ Therefore, private signals cannot have a large influence on an individual bidder's expectation.

3.4 Qualitative Results

In the empirical application of the procedure described in section 2.4 we set $\delta: = 0.01$, indicating an accuracy of the computation of bonds issued worth 10,000 Swiss francs. Convergence is achieved after eight iterations. The generated \tilde{B} -functions are decreasing for all 56 observations, so that we can say with some certainty that the constraint on the bidders to use downward-sloping bid functions imposed by NAUTZ [1995a] is empirically irrelevant for the Swiss Treasury bond auctions. As an illustration, figure 1 shows the process of convergence for a sample auction.

HARRIS AND RAVIV [1981, 1488] mention a set of empirical and experimental findings that relate to our hypothetical bid functions. They find:

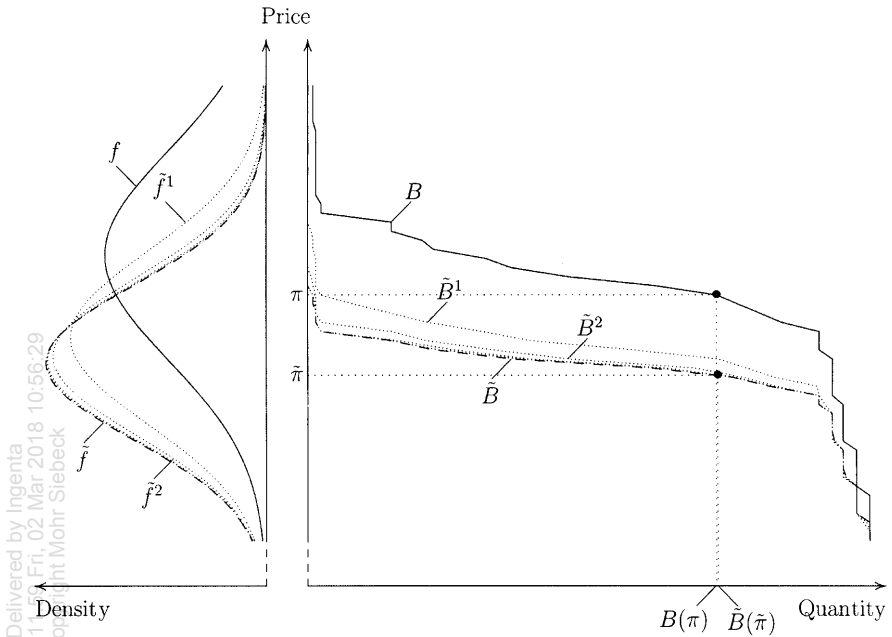
1. The mean bid is larger under the competitive than under the discriminating auction.
2. The variance of bids is larger under the competitive auction.
3. The evidence regarding the comparison of seller's revenue under the two types of auctions is inconclusive."

Point 1 is a corollary to proposition 1, and since our hypothetical bid functions satisfy this proposition by design, they also satisfy Harris and Raviv's first point. In fact, the hypothetical bid functions are strictly below the original bid functions for each of the 56 auctions, which implies that the mean bid is lower in each of the hypothetical price-discriminating auctions. Point 2 applies to our computed hypothetical bid functions because these functions always turn out to be flatter than the original bid functions in the uniform-price (or "competitive") auction. Thus, the variance of the bids under price discrimination, as we have computed them, is smaller than that of the bids under the uniform-price auction (figure 1 shows this clearly). The fact that the bid functions in the hypothetical price-discriminating auctions reflect these empirical findings should provide some confidence in the validity of the computation procedure used to generate those bid functions. Our

⁸ The regression also yields a constant (α) and a slope (β) that are statistically indistinguishable from zero and one, respectively, establishing that $\pi^s (= m)$ is an unbiased signal for π . More precisely, α is estimated to be -0.080 with standard error 1.331, and β to be 1.0004 with standard error 0.013.

Figure 1

Bid Functions (Right) and Density of Expected Cut-Off Price (Left) of a Sample Auction, Original Data, First Four Iterations, and Final Estimate



computation does not, however, produce the same inconclusive results indicated by point 3, i.e., we are able to rank the two auction forms with respect to their revenue implications. The next section discusses this aspect.

3.5 Quantitative Results

Table 1 contains all the results necessary for the comparison of the performance of the two auction rules. It features the price signal of the secondary market, π^s ; the price realized in the uniform-price auction, π ; the number of bonds issued, $B(\pi)$; and the amount of funds raised, $S(= \pi B(\pi))$.⁹ It also includes the cutoff price

⁹ $B(\pi)$ as in table 1 is simply the quantity of bids at or above the cutoff price. It can deviate from the number of bonds that were actually issued for two reasons. First, we disregard unpriced bids. We do so because our theory makes no prediction about why anyone would submit such bids. For this reason we have no description of such bidders' behavior in a price-discriminating auction. Second, we also disregard rationing of the lowest accepted class of bids, which occasionally occurred. Both omissions are of minor importance.

Table 1
Comparison of Actual and Simulated Auctions

Obs	π^s	π	$B(\pi)$	S	$\tilde{\pi}$	$\tilde{B}(\tilde{\pi})$	$\tilde{\pi} - \pi$	$\tilde{B}(\tilde{\pi}) - B(\pi)$	Advantage uniform
1	98.65	98.70	768.60	758.61	98.06	773.17	-0.64	+4.57	+0.59%
2	98.60	98.50	172.25	169.67	97.90	173.14	-0.60	+0.89	+0.52%
3	93.50	92.50	147.50	136.44	92.24	147.65	-0.26	+0.15	+0.10%
4	94.56	94.45	712.50	672.96	93.90	716.12	-0.55	+3.62	+0.51%
5	99.15	99.10	867.15	859.35	98.51	871.93	-0.59	+4.78	+0.55%
6	97.95	97.60	755.00	736.88	97.16	757.74	-0.44	+2.74	+0.36%
7	105.96	105.80	142.00	150.24	105.25	142.69	-0.55	+0.69	+0.49%
8	99.95	99.60	954.00	950.18	99.15	957.50	-0.45	+3.50	+0.37%
9	106.75	106.00	145.20	153.91	105.60	145.49	-0.40	+0.29	+0.20%
10	99.90	99.65	474.60	472.94	99.16	476.74	-0.50	+2.14	+0.45%
11	103.40	103.30	274.50	283.56	102.72	275.77	-0.58	+1.26	+0.46%
12	95.60	95.75	260.00	248.95	95.05	261.70	-0.70	+1.70	+0.65%
13	114.98	114.70	389.00	446.18	114.15	390.51	-0.55	+1.51	+0.39%
14	103.88	103.10	175.50	180.94	102.80	175.82	-0.30	+0.31	+0.18%
15	107.35	107.35	818.55	878.71	106.73	823.09	-0.62	+4.55	+0.56%
16	108.28	108.50	723.00	784.46	107.73	727.93	-0.77	+4.93	+0.68%
17	103.75	100.25	170.10	170.53	100.20	169.49	-0.05	-0.61	-0.36%
18	101.30	101.30	462.20	468.21	100.67	464.77	-0.63	+2.57	+0.56%
19	106.28	106.00	118.00	125.08	105.50	118.45	-0.50	+0.45	+0.38%
20	101.10	101.10	778.10	786.66	100.49	782.48	-0.61	+4.38	+0.56%
21	103.82	103.80	506.50	525.75	103.19	509.35	-0.61	+2.85	+0.56%
22	111.50	111.60	501.00	559.12	110.87	504.14	-0.73	+3.14	+0.63%
23	101.50	101.10	79.00	79.87	100.67	79.30	-0.43	+0.30	+0.38%
24	105.40	105.30	849.50	894.52	104.71	854.09	-0.59	+4.59	+0.54%
25	105.72	105.60	699.01	738.16	105.04	702.52	-0.56	+3.51	+0.50%
26	98.60	98.50	242.40	238.76	97.92	243.54	-0.58	+1.14	+0.47%
27	106.87	106.90	305.20	326.26	106.25	306.99	-0.65	+1.79	+0.59%
28	105.47	105.45	878.53	926.41	104.84	883.32	-0.61	+4.79	+0.55%
29	103.78	103.45	726.63	751.10	103.00	729.28	-0.45	+2.64	+0.36%
30	100.61	100.60	1079.50	1085.98	99.98	1085.53	-0.62	+6.03	+0.56%
31	111.73	111.50	294.70	328.59	110.96	295.93	-0.54	+1.23	+0.42%
32	100.65	100.35	452.00	453.58	99.86	453.76	-0.49	+1.76	+0.39%
33	103.40	103.10	300.60	309.92	102.63	301.81	-0.47	+1.21	+0.40%
34	104.53	104.50	219.00	228.86	103.88	220.24	-0.62	+1.24	+0.56%
35	99.52	100.75	68.02	68.53	99.25	69.05	-1.50	+1.03	+1.51%
36	105.29	105.40	302.55	318.89	104.70	304.54	-0.70	+1.99	+0.66%
37	92.75	92.75	93.00	86.26	92.08	93.50	-0.67	+0.50	+0.53%
38	97.80	97.90	585.20	572.91	97.24	589.20	-0.66	+4.00	+0.68%
39	96.75	96.80	516.15	499.63	96.17	519.50	-0.63	+3.35	+0.65%
40	100.72	100.90	605.12	610.56	100.18	609.53	-0.72	+4.41	+0.73%

Table 1
(Continued)

Obs	π^s	π	$B(\pi)$	S	$\tilde{\pi}$	$\tilde{B}(\tilde{\pi})$	$\tilde{\pi} - \pi$	$\tilde{B}(\tilde{\pi}) - B(\pi)$	Advantage uniform
41	101.46	101.50	505.07	512.65	100.86	508.04	-0.64	+2.97	+0.59%
42	101.55	101.50	429.54	435.98	100.91	431.79	-0.59	+2.25	+0.52%
43	100.08	100.00	347.32	347.32	99.41	349.25	-0.59	+1.93	+0.56%
44	93.00	92.40	194.50	179.72	92.06	194.85	-0.34	+0.34	+0.18%
45	98.85	98.40	200.18	196.98	98.00	200.82	-0.40	+0.64	+0.32%
46	100.20	100.30	299.14	300.03	99.60	301.03	-0.70	+1.90	+0.63%
47	95.35	95.10	351.80	334.56	94.60	353.05	-0.50	+1.25	+0.35%
48	99.90	99.80	257.50	256.99	99.23	258.79	-0.57	+1.29	+0.50%
49	97.90	97.30	337.75	328.63	96.95	338.30	-0.35	+0.55	+0.16%
50	97.96	98.40	197.42	194.26	97.50	199.10	-0.90	+1.68	+0.85%
51	99.15	99.30	484.70	481.31	98.60	487.76	-0.70	+3.06	+0.63%
52	98.80	99.20	303.60	301.17	98.34	306.20	-0.86	+2.60	+0.86%
53	101.60	101.20	423.52	428.60	100.76	424.18	-0.44	+0.66	+0.15%
54	102.65	103.00	895.67	922.54	102.15	902.82	-0.85	+7.15	+0.80%
55	102.75	103.10	511.30	527.15	102.25	515.44	-0.85	+4.14	+0.81%
56	102.70	103.10	368.81	380.24	102.22	371.80	-0.88	+2.99	+0.81%

computed for the hypothetical price-discriminating auction $\tilde{\pi}$ and the number of bonds that would have to have been issued in this auction in order to raise S funds, namely $\tilde{B}(\tilde{\pi})$. The next two columns show by how much the price and the quantity issued differ for the two formats. As one would expect, the discriminatory auction always has a lower cutoff price than the uniform-price auction. The last column reports by what proportion the Treasury has fared better using the uniform-price than using the discriminatory price auction format, i.e., the percentage by which $\tilde{B}(\tilde{\pi})$ exceeds $B(\pi)$.

The results demonstrate that, in the past, price discrimination would not have decreased the Treasury's cost. On the contrary, in all except one auction it would have required the Treasury to issue more bonds in order to collect the same amount of funds, so the uniform-price auction does have an advantage for the Treasury. Whether this difference is significantly larger than zero can be tested in various ways.

Since for every reopening we have one observation for each auction format, $\tilde{B}(\tilde{\pi})$ and $B(\pi)$ are two dependent samples. A common test for difference in such a case is a dependent-sample t -test. The null hypothesis under this test is that the average advantage of one of the auction formats is equal to zero. Given our observations for $\tilde{B}_z(\tilde{\pi}_z)$ and $B_z(\pi_z)$, we are able to reject the null hypothesis ($t = 10.41$).

Table 2
Testing for Cost Equivalence

Coefficient	Estimate	Std. error	<i>t</i> -statistic	Prob $\{H_0\}$
<i>a</i>	-0.2169	0.19770	-1.10	0.227
<i>b</i>	+0.0058	0.00039	15.02	0.000

Note: Adj. R^2 is 0.80.

Hence, the uniform auction is significantly better for the Treasury than the discriminating.

The use of the *t*-test presupposes the normal distribution of the differences between samples. Since it is unclear whether this assumption is valid, we can also check with a nonparametric test. The Wilcoxon matched-pairs signed rank test also leads to a rejection of the null hypothesis ($t = 6.43$) – this despite the lower power of nonparametric tests.

Furthermore, we can also consider the following regression:

$$(6) \quad \tilde{B}_z(\tilde{\pi}_z) - B_z(\pi_z) = a + b \cdot B_z(\pi_z) + \text{error}_z$$

where z denotes some auction. The left-hand variable is the number of bonds that must be issued in a price-discriminating auction in excess of what is issued in the uniform-price auction. The null hypothesis is that the uniform price and the discriminating price auctions are cost-equivalent, implying that both coefficients, a and b , are zero. This hypothesis can be rejected based on an OLS regression of (6); see table 2. The absolute term a is not significantly different from zero, but the slope coefficient b is clearly positive. The size of the slope implies that the price-discriminating auction would have increased the government's debt finance cost by 0.58%.

4 Conclusions

The question of which auction procedure the Treasury should use is the subject of ongoing debate. In this paper we have shown that theory alone cannot settle this question, nor can any of the empirical experiments carried out to date.

In an attempt to determine which auction format is better, some Treasuries have been running experiments using both auction methods. From an experimental design point of view, all of these past experiments are flawed, since differences in performance cannot be clearly attributed to changes in the auction format, either because the two auction formats are used at different points in time, or because they are used for different types of securities at the same point in time. Until now, researchers have been constrained to the use the data from such experiments to investigate this question.

In this paper, we develop a method for comparing the performance of the two auction formats even when no experiments are available. Our method allows us to generate the hypothetical bid functions that would have occurred under price discrimination using data that have been observed in uniform-price auctions only. Our method is therefore not subject to the identification problem inherent in the available experimental data.

Using our computation procedure, we conclude that the uniform-price auction has a statistically significant advantage over the price-discriminating auction in Switzerland. Compared to the price-discriminating auction, it reduces the government's debt financing cost by more than half a percentage point. Thus, the Swiss Treasury has wisely chosen to use the uniform-price format for all its auctions.

References

- BACK, K., AND J. F. ZENDER [1993], "Auctions of Divisible Goods: On the Rationale for the Treasury Experiment," *Review of Financial Studies*, 6, 733–764.
- AND – [1999], "Auctions of Divisible Goods with Endogenous Supply," Manuscript, Washington University in Saint Louis.
- BERG, S. A., B. BOUKAI, M. LANDSBERGER, AND Y. LENGWILER [2000], "Market Demand for Treasury Securities," Manuscript, Department of Economics, Haifa University, Israel.
- BIKCHANDARI, S., AND C.-F. HUANG [1993], "The Economics of Treasury Security Markets," *Journal of Economic Perspectives*, 7, 117–134.
- ENGELBRECHT-WIGGANS, R., AND C. M. KAHN [1998a], "Multi-Unit Auctions with Uniform Prices," *Economic Theory*, 12, 227–258.
- AND – [1998b], "Multi-Goods Pay-Your-Bid Auctions with Variable Awards," *Games and Economic Behavior*, 23, 25–42.
- FRIEDMAN, M. [1959], "Testimony in Employment, Growth, and Price Levels," Hearings before the Joint Economic Committee, 86th Congress, 1st Session, October 30.
- HARRIS, M., AND A. RAVIV [1981], "Allocation Mechanisms and the Design of Auctions," *Econometrica*, 49, 1477–1499.
- HORTACSU, A. [2000], "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market," Manuscript, Stanford University.
- LENGWILER, Y. [1998], "The Multiple Unit Auction with Variable Supply," *Economic Theory*, 14, 373–392.
- MALVEY, P. F., AND C. M. ARCHIBALD [1998], "Uniform-Price Auctions: Update of the Treasury Experiment," Technical Report, Office of Market Finance, U.S. Treasury.
- –, – –, AND S. T. FLYNN [1995], "Uniform-Price Auctions: Evaluation of the Treasury Experiment," Technical Report, Office of Market Finance, U.S. Treasury.
- MCADAMS, D. [2000], "Collusive-Seeming Equilibria' in the Uniform-Price Auction," Manuscript, Stanford University.
- NAUTZ, D. [1995a], "Optimal Bidding in Multi-Unit Auctions with Many Bidders," *Economics Letters*, 48, 301–306.
- – [1995b], "Zur Feinststeuerung des Geldmarktes durch die Wertpapierpensionsgeschäfte der Bundesbank," *Zeitschrift für Wirtschafts- und Sozialwissenschaften*, 115, 623–644.
- NOUSSAIR, C. [1995], "Equilibria in a Multi-Object Uniform Price Sealed Bid Auction with Multi-Unit Demands," *Economic Theory*, 5, 337–351.
- NYBORG, K. G. [1997], "On Implicit Collusion in Treasury Auctions," Manuscript, London Business School.

Delivered by Ingenta
Copyright © Siebert
11.591.01/Mar 2015 10:56:23

- —, AND S. SUNDARESAN [1996], “Discriminatory versus Uniform Treasury Auctions: Evidence from When-Issued Transactions,” *Journal of Financial Economics*, 42, 63–104.
- PAVAN, A., M. LICALZI, AND M. GILLI [1999], “A Simple Way to Inhibit Collusion in Uniform-Price Treasury Auctions,” Manuscript, University of Venice, Italy.
- SIMON, D. [1994], “The Treasury’s Experiment with Single-Price Auctions in the Mid 1970’s: Winner’s or Taxpayer’s Curse?” *Review of Economics and Statistics*, 76, 754–760.
- SMITH, V.L. [1966], “Bidding Theory and the Treasury Bill Auction: Does Price Discrimination Increase the Bill Price?” *Review of Economics and Statistics*, 48, 141–146.
- TENORIO, R. [1993], “Revenue Equivalence and Bidding Behavior in a Multi-Unit Auction Market: An Empirical Analysis,” *Review of Economics and Statistics*, 75, 302–314.
- — [1997], “On Strategic Quantity Bidding in Multiple Unit Auctions,” *Journal of Industrial Economics*, 45, 207–217.
- UMLAUF, S.R. [1993], “An Empirical Study of the Mexican Treasury Bill Auction,” *Journal of Financial Economics*, 33, 313–340.
- WANG, J.J.D., AND J.F. ZENDER [1999], “Auctioning Divisible Goods,” Manuscript, Fuqua School of Business, Duke University.
- WILSON, R. [1979], “Auctions of Shares,” *Quarterly Journal of Economics*, 93, 675–689.

Daniel Heller
Yvan Lengwiler
Swiss National Bank
P.O. Box
8022 Zürich
Switzerland
E-mail:
daniel.heller@snb.ch
yvan.lengwiler@snb.ch