Initiative for Policy Dialogue Working Paper Series

September 2006

How Best to Auction Oil Rights

Peter Cramton

Resource Curse

Chapter 5: How Best to Auction Oil Rights?

Peter Cramton[*]

Abstract

A good auction design promotes both an efficient assignment of rights and competitive revenues for the seller. The two key factors that determine the best design are the structure of bidder preferences and the degree of competition. With weak competition and "additive values," a simultaneous first-price sealed-bid auction may suffice. With more complex value structures, a dynamic auction with package bids, such as the clock-proxy auction, is likely needed to increase efficiency and maximize revenues. Bidding on production shares, rather than bonuses, typically increases Government Take by reducing oil company risk.

Introduction

There are many ways for states to assign oil rights. Rights are sometimes assigned via informal processes, such as first-come-first-serve, or other processes, such as "beauty contests," in which companies submit exploration and development plans. In this chapter, however, I examine the design of auctions for oil rights, focusing especially on issues faced in developing countries. For the purposes of this chapter, I assume that revenue maximization is the overriding objective. While certainly there can be other objectives, such as the timing of the revenues and country employment and investment (see Chapter 3), in what follows, I assume that revenue is the main objective. In this context, the advantage of an auction is that it is a transparent method of assignment, which if well designed, is capable of maximizing the revenues that a developing country can receive from its oil and gas endowments. Nonetheless there are advantages and disadvantages of different auction designs for oil and gas producing developing countries, which renders the design of auctions especially important.

Design and process issues are in fact especially important with developing countries. While it is always necessary to tailor a design to a particular setting, there are a number of general insights we can draw from recent auction theory and practice, both in oil rights auctions and in other sectors. The first step is defining the product: the term of the license, the block size, royalties and tax obligations, and so on. Next, a number of basic design issues must be resolved: whether to sell rights sequentially or simultaneously, whether to use a dynamic or a static auction, what information policy should be used, and whether and how reserve prices should be set. In considering these questions, risks of collusion and corruption must also be examined.

Much depends on the structure of bidder preferences (or 'values'). Two aspects of bidder preferences are especially important. The first is that the values to a bidder of a particular item might depend on what other items he already owns. Items for sale -- the right to explore and develop oil and gas on a particular geographic block -- are sometimes "substitutes" and sometimes "complements," in the sense that sometimes possessing one right (a block for instance) makes other rights less valuable for a particular bidder, but sometimes possessing one right makes other rights more valuable. If for a particular bidder, the value of one block is independent of ownership of another, then we say that the values are additive. The second is that unlike many settings, the *values* that the bidders place on oil and gas rights may be interdependent across the different bidders, since each bidder has private information, from surveys and seismic tests, which is relevant in determining the overall value of a block. Bidders have "*common* values" if it is the case that the ex post value of the block is the same for all bidders. This ex post value is unknown at the time of the auction; bidders only have estimates of the common value from the surveys, seismic tests, and expert analysis they have conducted.

Auction theory suggests that when bidders have preference structures like this -- viewing blocks as substitutes or as complements and having common values with private information, then some version of a *simultaneous ascending auction* is best, since this will promote efficient pricing and packaging of the blocks. In brief, a simultaneous ascending auction is characterized by the simultaneous sale of many related blocks using an ascending price process in which bidders can improve their bids for the blocks. The auction ends when no bidder is willing to bid higher for any block.

There are a number of different types of simultaneous ascending auctions.

At one extreme is the clock-proxy auction (Ausubel, Cramton, and Milgrom 2006), a sophisticated version of the simultaneous ascending auction often used in the auction of radio spectrum. The mechanics of the clock-proxy auction are discussed in greater detail below. The clock-proxy auction is a method of auctioning many related items over multiple bidding rounds, allowing bids on packages of items. The auction begins with a clock phase. The auctioneer names a price for each block and the bidders respond with the set of blocks they desire at the specified prices. Prices increase on blocks with more than one bid. This process continues until there are no blocks with more than one bid. After this, there is a final proxy round in which bidders express values for any desired packages of blocks. An efficient assignment of blocks is found based on the proxy bids and all the bids in the clock phase; prices are determined, and an efficient assignment of blocks is found.

The clock-proxy auction encourages effective price discovery in the clock phase and the proxy round promotes an efficient assignment and competitive revenues. Although this approach may appear complex, it is actually simpler for bidders than common alternatives. The reason is that, although the rules may appear complex, the best strategies for bidders are simple. The price discovery reduces guesswork and focuses the bidders' attention on the relevant part of the price space; then the proxy round gives the bidders a means to further express package preferences and fine-tune the assignment of blocks. The approach is well suited for excellent prospects, with complex value structures, like those described above.

At the other extreme is the first-price sealed-bid auction used in the U.S. for offshore leases. The bidders simultaneously submit bids for each desired block. Each block is awarded to the highest bidder at the winning bid price. This simple format is suitable for marginal blocks with nearly *additive* value structures (the value of a package is equal to the sum of the values of the individual blocks) and *small value interdependencies* across bidders.

Still other designs between these two extremes are appropriate when the bidder preferences are not so complex that package bidding is essential and not so simple as additive values. In the remainder of this chapter, I will develop the logic of these different types of auction in more detail and discuss when one type of design is likely to work better than another. I begin with some motivating insights from auction theory and practice. I then consider bidder preferences and some of the basic design issues in oil rights auctions. The following three sections address problems specific to developing countries and the experience with oil rights auctions and auctions in other sectors. I then describe the clock-proxy auction and in the final section consider a number of alternative auction formats and make recommendations based on the particular setting.

Motivating insights from auction theory and practice

Why auction?

Auctions allocate and price scarce resources in settings of uncertainty. Every auction asks and answers the basic question: who should get the items and at what prices? Auctions are a formal and transparent method of assignment. Clear rules are established for the auction process. Transparency benefits both the bidders and the country. It mitigates potential corruption and encourages competition through a fair and open process.

A primary advantage of an auction is its tendency to assign the blocks to those best able to use them. Although this does not always occur (a number of features that limit the efficiency of auctions are discussed below), the competitive character of auctions makes it more likely: Companies with the highest estimates of value for the blocks are likely to be willing to bid higher than the others, and hence tend to win the blocks.

Informal processes, such as negotiation on a first-come-first-serve basis, lack transparency and are vulnerable to favoritism and corruption, which undermines competition. The reduced competition inherent in an informal process reduces both the efficiency of the assignment and the country's revenues. Informal processes also tend to be more vulnerable to expropriation, further discouraging competition.

A leading alternative to auctions is an administrative process, often called 'beauty contests,' in which oil companies present plans for exploration and development according to a formal process. This approach may be more flexible than auctions, but it makes the assignment less transparent and more vulnerable to favoritism and corruption.

Does auction design matter?

One of the most important results of auction theory is the "revenue equivalence theorem." The revenue equivalence theorem makes a remarkable claim: under particular assumptions, the four standard methods for auctioning a single item (first-price sealed-bid, second-price sealed-bid, English ascending, and Dutch descending)[1] all result in exactly the same expected revenue for the seller. In each case, the expected revenue is equal to the expected value placed on the item by the bidder who values it the second most. Furthermore, when the seller sets an appropriate reserve price, these four methods all result in revenues that are at least as large as those that can be achieved from any other trading mechanism. From this striking result, one might conclude that auction design is of little importance -- that all standard auctions perform well.

This, however, is the wrong conclusion. The assumptions required for the revenue equivalence theorem are quite special; notably the theory assumes that sellers are auctioning a single item, that bidders have independent private values, that bidders are risk neutral, that the number of bidders is independent of the type of auction used and that there is no collusion or corruption, and except for their different valuations of the good, bidders are otherwise identical to each other. In practice, none of these assumptions holds: many related items are for sale; bidder values depend at least in part on value estimates of other bidders and these estimates are correlated; bidder participation decisions are of paramount importance; bidders care about risk; there are ex ante differences among the bidders (e.g., some are large and some are small); and mitigating collusion and corruption are important. Each of these features impacts the performance of alternative auction designs. The choice of the best auction design depends on which of these different features are most salient.

Product definition

The first step is product definition -- what is being sold. There are two key elements: 1) the contract terms of the license (duration, royalties, tax obligations), and 2) the geographic scope of the blocks. The first of these is discussed in more detail in Chapter 3. The determination of the second depends to a large extent on local context and I do not discuss it at length here. In general, however, blocks are defined as rectangular blocks, as specified by a pair of longitude and latitude coordinates. The appropriate size of the blocks depends on the quality of the prospect. More promising regions support smaller blocks. In the U.S., blocks are nominated by the oil companies. This is a sensible approach in most cases because it guarantees at least some interest in the auctioned blocks.

Auction process

As important as the auction design itself is, the process through which the auction takes place is equally important. To promote transparency, the auction process must be specified well in advance of the tender. The process should be open to all oil companies on a nondiscriminatory basis. The process should begin with a public advertisement of the tender and a complete description of the procedure for awarding a license should be provided, including bidder qualification procedures and the auction rules. Such a clear and complete statement of the auction process is essential to bidder participation. The country should be committed to the process. Finally, the process should allow for and encourage input from the oil companies in a transparent setting with "sunshine" rules that require public announcement of the content of meetings between the country and the oil companies. At a minimum, this would include the nomination of blocks, but allowing comments on all aspects of the rule making is generally worthwhile. Bidder participation and bids are enhanced if legitimate bidder concerns and preferences are addressed.

Bidder preferences and auction design

The structure of bidder preferences

Before considering design issues, it is helpful to think first about the bidders' preferences. We consider two aspects of bidder preferences that affect the optimal design of auctions: the interdependence of valuations across bidders and the interdependence of valuations across blocks.

Interdependence of Valuations Across Bidders. In the study of auctions there are three standard models for describing the valuations of bidders: private values, common values, and interdependent values.

If there are *private values* then this means that each bidder's value does not depend on the private information of the other bidders. Each bidder has its own valuation of the expected worth (to it) of the different packages of items on sale.

If there are *common values*, then packages of items have the same value to all bidders. But these values are unknown. The value can be written as a function of the individual bidder's private information, as well as the information held by all other bidders. The more a given bidder knows about other bidder's valuations, the better is that bidder able to estimate the worth of the item.

If there are *interdependent values*, then each bidder's value of a package depends on his private information as well as the private information of the other bidders. This is a more general formulation and both private and common values can be written as special cases of interdependent values. With interdependent values, each bidder has its own estimates of the value which is a function of the bidders own information and may also be a function of the other bidders' information. In this situation, a bidder's value depends on the values assigned by other bidders.

The oil rights setting is the textbook example of a common values auction. All companies value the oil at about the same amount (the world price of oil), but there is enormous uncertainty about the quantity of oil and the cost of extracting it. Before bidding, each company estimates these uncertainties from geological surveys, seismic tests, and analysis of petroleum engineers. Yet each company would like to have the private information of the other bidders to further reduce uncertainty. The common value depends not just on the bidder's estimate of value, but on all the other estimates. In practice, there are also some private value elements -- the company's exploration and development capacity, its reserves, its expertise in the particular type of prospect, its ability to manage exploration and political risks -- but these elements typically are of secondary importance. Thus, the oil rights setting has interdependent values with strong common value elements.

In situations in which there are common values, the strategies chosen by bidders are conditioned by a phenomenon known as the *winner's curse*. This is the insight that winning an item in an auction is actually bad news for the winner about the item's true value, because winning implies that no other bidder was willing to bid as much for the item. Hence, it is likely that the winner's estimate of value is an overestimate. Since a bidder's bid is only relevant in the event that the bidder wins, the bidder should condition the bid on the negative information winning conveys about value. Bidders that fail to condition their bids on the bad news winning conveys suffer from the winner's curse in the sense that they often pay more for an item than it is worth. In oil rights auctions, adjusting bids in light of the winner's curse is a key element of strategy. In contrast, in private values auctions, there is no winner's curse: each bidder knows what value they place on the object and this value does not depend on the values of the others.

Interdependence of Valuations Across Blocks. Thus far, we have focused on how package values depend on private information. A second important dimension is the structure of *package* values. How does the bidder value a package of blocks?
The simplest valuation model is *additive values*: the value of a package is the sum of the values of the individual blocks. In oil rights auctions, additive values are a good first approximation. The primary determinant of value is the quantity of oil, and the quantity of oil in a package of blocks is simply the *sum* of the quantities in each block.

However, sometimes values may also be either *sub*additive or *super*additive.

With *subadditive values*, the value of a package is less than the sum of the individual values. One source of subadditive values is capacity constraints on exploration and refining. Additional blocks have less value if the company lacks the resources to efficiently exploit that value. Another source is risk, holding many blocks within the same region where values are highly correlated is riskier than holding a few blocks in each of many dispersed regions. Values for substitute goods are subadditive.

With *superadditive values*, the value of a package is greater than the sum of the individual values. Superadditive values arise if there are synergies or blocks act as complements. One source of complements is exploration and production efficiencies that arise from holding many neighboring blocks. Traditional economies of scale may arise in drilling from sharing staff and equipment. A more subtle form of complements comes from more efficient exploration. For example, if two neighboring blocks are owned by different companies, each may have an incentive to free ride on the exploration efforts of the other -- waiting to see if the other's drilling is successful. As a result, the exploration of both tracks may be inefficiently delayed. Hendricks and Porter (1996) provide both a theoretical model and empirical support for this behavior in the U.S. offshore oil lease auctions. If instead, the two blocks are held by the same company, there is no information externality and the blocks are explored efficiently. A related synergy comes from the common pool problem, in which neighboring blocks are drawing oil from the same pool. When the blocks are held by the same company, the exploitation of the pool is efficient; whereas, with separately held blocks, the companies would need to negotiate a unitization agreement to coordinate the development. Ideally, blocks are defined to avoid this problem, but the country may not have sufficient information to avoid it entirely.

In the oil rights setting, additive values may be a good first approximation. Nonetheless, complements (superadditivity) and substitutes (subadditivity) likely are important in at least some applications. If this is the case, then the auction design needs to allow for efficient packaging. Otherwise, if values are largely additive, then packaging issues can be safely ignored, resulting in a much simpler auction design.

Basic design issues

Given this characterization of the different ways that bidders might value a set of blocks I now address several key issues of auction design in the oil rights setting.

Open bidding or sealed bidding?

Especially in contexts where there is interdependence of valuations across bidders, open bidding is generally better than a single sealed bid. An essential advantage of open bidding is that the bidding process reveals information about individual valuations. This information promotes the efficient assignment of licenses, since bidders can condition their bids on more information. Moreover, since bidders' private information is likely to be positively correlated, open bidding may raise auction revenues (Milgrom and Weber 1982). Intuitively, bidders are able to bid more aggressively in an open auction, since they have better information about the item's value. The open bidding reveals information about the other bidders' estimates of value. This information reduces the bidder's uncertainty about value, and thus mitigates the winner's curse -- the possibility of paying more than the value of the item. Thus, bidders are able to bid more aggressively, and this translates into high revenues for the seller. That turns out to be a strong argument in favor of open rather than sealed bidding.

Sealed bidding has some advantages. Most importantly, a sealed-bid design is less susceptible to collusion, such as agreements among oil companies not to compete against each other (Milgrom 1987). Open bidding allows bidders to signal through their bids and establish tacit agreements. With open bidding, these tacit agreements can be enforced, since a bidder can immediately punish another that has deviated from the collusive agreement. Signaling and punishments are not possible with a single sealed bid. So in situations in which collusion is a major concern sealed bidding may make more sense.

A second advantage of sealed bidding is that *if there are ex ante differences among the bidders*, it may yield higher revenues (Maskin and Riley 2000, Klemperer 2002). This is especially the case if the bidders are *risk averse* and have *independent private values*. The reason is that in a sealed-bid auction, a strong bidder can guarantee victory only by placing a high bid; even if a strong bidder that expects to win has a good sense of what the next most likely winner is willing to bid, they are likely to bid some margin above this in order to avoid risking losing the block. In an open auction on the other hand, the strong bidder never needs to bid higher than the second-highest value.

In the oil rights auctions, an open auction probably is best, provided the design adequately addresses potential collusion. The reason is that values have a strong common value element and so the benefits of sealed bids are not so great but the benefits of an open auction are substantial.[2]

Should auctions of multiple blocks be run simultaneously or sequentially?

Generally, simultaneous open bidding is better for the seller than sequential auctions. A disadvantage of sequential auctions is that they limit the information available to bidders and limit how the bidders can respond to information. With sequential auctions, bidders must guess what prices will be in future auctions when determining bids in the current auction. Incorrect guesses may result in an inefficient assignment when item values are interdependent. A sequential auction also eliminates many strategies. A bidder cannot switch back to an earlier item if prices go too high in a later auction. Bidders are likely to regret having purchased early at high prices, or not having purchased early at low prices. The guesswork about future auction outcomes makes strategies in sequential auctions complex, and the outcomes less efficient.

In a simultaneous ascending auction, a large collection of related items is up for auction at the same time. Hence, the bidders get information about prices on all the items as the auction proceeds. Bidders can switch among items based on this information. Hence, there is less of a need to anticipate where prices are likely to go. Moreover, the auction generates market prices. Similar items sell for similar prices. Bidders do not regret having bought too early or too late and with less fear of such regrets are willing to bid more.

Proponents of sequential auctions argue that the relevant information for the bidders is the final prices and assignments. They argue that simultaneous auctions do not reveal final outcomes until the auction is over. In contrast, the sequential auction gives final information about prices and assignments for all prior auctions. This final information may be more useful to bidders than the preliminary information revealed in a simultaneous auction. However, empirical analysis of simultaneous ascending auctions indicates that the preliminary information is highly correlated with final outcomes (Cramton 1997).

Supporters of sequential auctions also point out that the great flexibility of a simultaneous auction makes it more susceptible to collusive strategies. Since nothing is assigned until the end in a simultaneous auction, bidders can punish aggressive bidding by raising the bids on those items desired by the aggressive bidder. In a sequential auction, collusion is more difficult. A bidder that is supposed to win a later item at a low price is vulnerable to competition from another that won an earlier item at a low price. The early winner no longer has an incentive to hold back in the later auctions. This potential problem, however, can be mitigated through an effective information policy, which determines what information is made public during the auction.

In oil rights auctions, the virtues of the simultaneous auction -- greater information release and greater bidder flexibility in responding to information -- would improve efficiency. But as with our arguments for adopting open bidding over sealed bidding, this depends in part on how well concerns about bidder collusion can be addressed.

Should bidders bid for individual blocks separately or for packages of blocks?

In general, it is a good idea for the seller to allow bidders to give different bids for different packages. Package bidding is desirable when a bidder's value of a block depends on what other blocks it wins. Package bidding also has advantages when bidders have budget constraints or other constraints that depend on the package of blocks won. Then bidders may prefer being able to bid on a combination of blocks, rather than having to place a number of individual bids. With a package bid, the bidder either gets the entire combination or nothing. There is no possibility that the bidder will end up winning just some of what it needs or that it wins more than it wishes to pay for.

With individual bids, bidding for a synergistic combination is risky. The bidder may fail to acquire key pieces of the desired combination, but may pay prices based on the synergistic gain. Alternatively, the bidder may be forced to bid beyond its valuation in order to secure the synergies and reduce its loss from being stuck with some low-value blocks. This is called the *exposure problem*. Individual bidding exposes bidders seeking synergistic combinations to aggregation risk.

To see how not allowing package bids can create inefficiencies, consider the following example of two bidders bidding for two adjacent parking spaces. One bidder with a car and a trailer requires both spaces. She values the two spots together at $100 and a single spot is worth nothing; the spots are perfect complements. The second bidder has a car, but no trailer. Either spot is worth $75, as is the combination; the spots are perfect substitutes. Note that the efficient outcome is for the first bidder to get both spots for a social gain of $100, rather than $75 if the second bidder gets a spot. Yet any attempt by the first bidder to win the spaces is foolhardy. The first bidder would have to pay at least $150 for the spaces, since the second bidder will bid up to $75 for either one. Alternatively, if the first bidder drops out early, she will "win" one license, losing an amount equal to her highest bid. The only equilibrium is for the second bidder to win a single spot by placing the minimum bid. The outcome is inefficient, and fails to generate *any* revenue. In contrast, if package bids are allowed, then the outcome is efficient. The first bidder wins both spots with a bid of $75 for the pair of spots.

This example is somewhat extreme but it illustrates well the exposure problem. The inefficiency involves large bidder-specific complementarities and a lack of competition. In practice, the complementarities in oil rights auctions are likely to be less extreme and the competition is likely to be greater.

Unfortunately, allowing package bids creates other problems. Package bids may favor bidders seeking large aggregations due to a variant of the free-rider problem, called the *threshold problem*. Continuing with the last example, suppose that there is a third bidder who values either spot at $40. Then the efficient outcome is for the individual bidders to win both spots for a social gain of 75 + 40 = $115. But this outcome may not occur when values are privately known. Suppose that the second and third bidders have placed individual bids of $35 on the two licenses, but these bids are topped by a package bid of $90 from the first bidder. Each bidder hopes that the other will bid higher to top the package bid. The second bidder has an incentive to understate his willingness to push the bidding higher. He may refrain from bidding, counting on the third bidder to break the threshold of $90. Since the third bidder cannot come through, the auction ends with the first bidder winning both spaces for $90.

A second problem with allowing package bids is complexity. If all combinations are allowed, even identifying the revenue maximizing assignment is a difficult integer programming problem when there are many bidders and items. Nonetheless, our understanding of and experience with package auctions has advanced considerably in recent years (Cramton, et al 2006). Package bids should therefore now be considered to be a viable option. Whether package bids are the best option will depend somewhat on the details of the setting.

How should reserve prices be used?

Reserve prices in oil rights auctions have two main purposes: 1) to guarantee substantial revenue in auctions where competition is weak but the reserve is met, and 2) to limit the incentive for -- and the impact of -- collusive bidding. Reserve prices mitigate collusive bidding by reducing the maximum gain of the collusive bidding. In effect this heightens the competition between bidders by increasing the importance of the gains they can make unilaterally relative to the gains they can make through collusion. Setting reserve prices for oil rights auctions is difficult given the enormous uncertainty of values. There is, however, an alternative to posting a reserve price that can be considered and that has for example been used in the U.S. It is possible for the seller to place a low minimum bid that applies to all blocks, and then accept or reject winning bids *ex post* (note the idea is to decide ex post whether to accept or reject a *winning* bid not to pick and choose between bidders ex post). Under such a system a reserve price exists but it is secret and can in fact depend on the observed bidding behavior.

Bonus bid, royalties, and production sharing

Oil rights auctions commonly involve bonus bids and either royalties or production sharing (see Chapters 3 and 2 for further discussions of the merits of these different approaches). The bonus bid or signature bonus is the upfront payment determined in auction for the right to explore and develop the block during the license period. If exploitable reserves are found, the license is renewed for a nominal fee as long as development continues. The royalty is the share of the oil and gas revenues that goes to the government. Royalty rates vary country to country and even within countries. For example, in the U.S. offshore oil lease auctions, the royalty rate is 1/6, where one-sixth of revenues for any oil extracted is paid as a royalty; whereas, the onshore rate typically is 1/8. The motivation for royalties is to have the oil company payment more closely reflect ex post realized value. This reduces the risk of the oil company. The disadvantage of royalties is that, like a tax, it distorts investment decisions. A larger royalty rate reduces the incentive for the oil company to invest in exploration and development activities. In contrast, the signature bonus is a sunk cost after the auction and does not distort subsequent investments. In a setting where there is no uncertainty about values, then only a bonus bid is needed (a zero royalty rate); in a setting where exploration and development are costless, then a 100 percent royalty rate is optimal. In practice, oil rights auctions have large uncertainty about values as well as large exploration and development costs. Thus, an intermediate rate is generally best.

Production sharing contracts attempt to further reduce oil company risk and better manage investment incentives by specifying the terms of cost sharing and profit sharing throughout exploration and development. The contract can allow the oil company to recover exploration and development capital costs before the country shares in the revenues. Then the government's profit share increases with the success of the project, allowing the terms to handle both marginal and windfall economics. The contracts often are made immune to tax changes by having the government counterparty, typically the national oil company, liable for all taxes. Work programs specify a minimum level of exploration effort. This is an important constraint on more marginal blocks, where high government profit shares might otherwise discourage exploration.

With production sharing contracts, it is common for bidding to be over the government's highest profit share, rather than the signature bonus. Thus, bidders compete on their willingness to share profits in the most favorable circumstances. This approach, used recently in Libya and Venezuela, reduces oil company risk without upsetting development incentives, since the bid share only applies for blocks that are highly successful. Development incentives are further maintained by having the government share in the development capital costs and the operating costs. If the government's share of development capital and operating costs is the same as its production share, then post-exploration the project essentially is a joint venture with first-best incentives for development.

Problems specific to developing countries

Developing countries face additional challenges in establishing an effective auction program. These include political risk, risk for companies of expropriation, favoritism, and corruption. All of these challenges tend to discourage participation in auctions and so reduce competition. In fact the strongest indicator of success of the auction program is the presence of robust competition. The geological prospect of the region is a primary factor in attracting oil companies, but political, legal, and process factors are also key.

There is little a country can do in the short term to reduce perceptions of political risk. Legal risks can be reduced through choice of contract law. And over the medium term, institutions can be developed that provide the ground rules for oil exploration and development.

Companies' fear of expropriation or adverse renegotiation can be mitigated somewhat through the cash flow structure of the contract terms. For example, a pure bonus bid system (zero royalty) is problematic in light of expropriation risks. This would force the oil company to sink most funds upfront, making the company vulnerable to expropriation. Even developed countries, such as the U.K. and the U.S., have a tendency to adjust tax rates after companies have begun production to capture a larger share of "windfall" profits. As a result, companies heavily discount bonus bids. Some reliance on royalties or production sharing is better, since these payments are not due until after revenues or profits have been received by the oil company. Another option is share bidding in which oil companies offer equity shares in the venture (the highest offered share wins the block). In this case, the country and the oil company are partners. Each makes investments and reaps rewards according to its share. This approach shifts risks from the oil companies to the country. More importantly, it aligns the interests of the company and the country, reducing expropriation risks for the company.

Such approaches are mutually beneficial but sometimes appear unattractive for developing countries. Developing countries, especially small ones, may have important constraints with respect to cash flows. For example, a country may be unable to make upfront outlays and so have strong preferences for early payments. However, too much focus on early revenues may greatly reduce total revenues, especially in an environment where renegotiation risk is high. For this reason, in the medium to long run, countries often are better off with production sharing contracts with small upfront payments and large government shares in the event of successful finds.

Finally, favoritism and corruption can be addressed in the auction process. A transparent, nondiscriminatory process is the key to mitigating favoritism and corruption. Independent third-party auction managers can help as well. Likewise, a trustee observing and commenting on all aspects of the auction process can further reduce the possibility of corruption. This approach is commonly used in developed countries. For example, electricity auctions in restructured electricity markets in North America and Europe are typically conducted by independent third-parties, often with a further independent review and certification of the entire process by an auditor or trustee.

Experience with oil rights auctions

Oil rights have been auctioned in many countries throughout the world. Much can be learned from these experiences. Here, I focus on the experiences of a wealthy nation, the United States, and compare this with recent experiences elsewhere, notably in Venezuela and Libya.

The U.S. experience

The most studied program is the U.S. offshore oil lease auctions. The discussion that follows draws largely on Porter's excellent survey of this research (Porter 1995). These oil lease auctions began in 1954. The product auctioned is a lease granting the right to explore and develop a particular tract for a period of five years (U.S. auctions use the terms 'lease' and 'tract', rather than 'license' and 'block'). If oil is found and developed, the lease is renewed for a nominal fee as long as production continues. The process begins with the oil companies nominating tracts for auction. The government then makes a list of tracts to be auctioned. The auction is a simultaneous first-price sealed-bid auction. Each bidder simultaneously submits a bid on each of the tracts it desires. The bid must meet or exceed the minimum bid, which is stated as a dollar amount per acre. The per-acre minimum depends only on the type of tract. A tract is either awarded to the high bidder or all bids on the tract are rejected; thus, the reserve price is secret and determined after the bids are observed by the government. A winning bidder pays its bid, which is referred to as the *bonus*. In addition, the company pays a royalty of 1/6 of revenues for any oil extracted. Bidders are allowed to bid jointly; however, after 1975, none of the top-eight oil companies could combine in a joint bid with another top-eight company.

Tracts are of three types. Wildcat tracts are new offerings that are not adjacent to developed tracts; drainage tracts are adjacent to developed tracts; and development tracts are a reoffering. There is an important economic difference between wildcat tracts and drainage tracts and these differences are reflected in the types of auctions used. With a drainage tract, bidders holding leases on adjacent tracts may have a much better estimate of value than those without adjacent tracts. Thus, the drainage tract sales may have large asymmetries among the bidders; whereas in the wildcat sales bidders are more symmetric. This difference has important implications for both bidding behavior and auction design. In particular, one would expect that a simultaneous ascending auction would be best for the wildcat tracts and a simultaneous sealed-bid auction would be best for the drainage tracts. However, the United States, perhaps for reasons of inertia, uses the same sealed-bid auction method for all tracts.

From 1954 to 1990, there were 98 auctions. On average, 125 leases sold per auction. Eight percent of the high bids were rejected. The auctions raised $55.8 billion from bonus bids and $40 billion from royalties (1972 dollars). Hendricks, Porter, and Boudreau (1987) estimate from ex post price and quantity data that the government share of value (revenue net of costs) was 77 percent with the oil companies receiving the remaining 23 percent.

Porter (1995) concludes that the U.S. auction program in many respects is well designed. Certainly the government is getting the lion's share of the value. On drainage tracts, informed bidders (those with leases on adjacent tracts), reap informational rents. The government could consider using a higher royalty rate on these tracts to the extent that the informational rents are not capitalized in the earlier wildcat sales.

One potentially troubling feature of the U.S. offshore program is the use of the simultaneous first-price sealed-bid format. This is easy for the government to implement, but poses challenges to bidders, which may reduce efficiency and revenues. In particular, the format prevents the bidders from expressing preferences for packages of tracts and it provides no price discovery. In addition, a bidder's budget constraints or other package-based constraints either cannot be satisfied or can only be satisfied by greatly distorting one's bids.

Onshore auctions in the U.S. are conducted at the state level. These auctions often are done as sequential open outcry auctions: each tract is sold in sequence using an English auction.

This approach allows for some price discovery and better handles budget constraints, but it still forces bidders to guess auction prices for leases sold later.

Experience outside the U.S.

Unfortunately, there is little publicly available information about oil rights auctions in developing countries, and little research on the topic (although one important exception is Sunley, Baunsgaard and Simard (2002) who provide a study of government revenue sources from oil and gas in developing countries). Typically however, countries employ a range of revenue methods: bonus bids, royalties, production sharing, income taxes, and state equity. Not surprisingly, the terms vary widely across countries, reflecting at least in part differences in political risks and geological uncertainty. A reasonable conclusion from the experiences of these countries is that auctions are a desirable method of allocating the rights among companies, but multiple revenue sources should be used to best manage risks and incentives.

Recent auctions conducted in an environment of high oil prices have been highly competitive, especially in regions with known reserves. For example, in the Libyan auction of 15 blocks on 29 January 2005, some blocks received as many as 15 bids.

Johnston (2005) examines the contract terms and bidding in the 2005 Libyan auction. This case study offers insights into modern contract terms and bidder competition in a major auction of excellent prospects during a period of high price expectations. The 15 blocks were offered in a simultaneous sealed-bid auction, in which oil companies bid a production share and a signature bonus for each desired block. Each license was awarded to the company with the highest production share (share of gross revenues going to the government). In the event of a tie, the signature bonus was used as a tie breaker.

The contract terms fully specify the split of revenues and costs between the government and the oil company. Companies bid only on the production share and the signature bonuses. Hence, for example, on block 54, the winning bid offered a production share of 87.6 percent. In effect this means that that the government gets 87.6 percent of the gross revenues, for which it pays none of the exploration costs, 50 percent of the development capital, and 87.6 percent of the operating costs. The oil company uses the remaining 12.4 percent of the gross revenues to recover its costs (100 percent of exploration costs, 50 percent of development capital, and 12.4 percent of operating costs). Once these costs are recovered from the 12.4 percent, the excess ("profit oil") is split between the government and the oil company according to a sliding scale based on a revenue/cost index. The government's share of this excess increases from 10 percent to 50 percent as the company's revenue/cost index increases from 1.5 to 3. Under these terms, the initial upfront capital expense is limited to the exploration cost and a modest signature bonus. Since development capital costs are split 50-50, the high production share does mean that some profitable fields may go undeveloped. However, once development capital is sunk, the 87.6-12.4 split of operating costs results in first-best incentives for extraction.

Competition in the Libyan round was intense with an average of 7 bidders per block. The winning production shares ranged from 61.1 to 89.2 percent with a mean of 80.5 percent. The Government Take (share of project profits) depends on the assumptions one makes on costs and revenues. Johnston (2005) estimates the Government Take to range from 77.0 to 97.7 percent with a mean of 89.9 percent, well above the 80 percent that is more typically captured for good prospects or the 77 percent realized in the U.S. auctions before 1990.

The 1996 Venezuela auction of 10 blocks had similar contract terms and also was highly successful. There were however some important differences. The 10 blocks were offered in sequence. Also, to maintain better development incentives, the production share bids were capped at 50 percent. First, the bidders bid production shares, and then in the event of a tie (e.g., two or more bid 50 percent) the bidders bid signature bonuses to break the tie. This resulted in large signature bonuses for desirable blocks, shifting risk to the winning oil companies. However, the Venezuela terms were more favorable than the Libya terms with respect to cost recovery, so it is unclear which terms were riskier. Indeed, the Government Take estimate of 92 percent remains a landmark figure (Johnston 2005).[3]

Both of these examples illustrate the success of first-price sealed-bid auctions. In light of the revenue equivalence theorem, this is not surprising. All competitive auctions, regardless of design, should generate substantial government revenues for excellent prospects. However, these examples do not show that the chosen approach was best. Indeed, revenues would be even higher if a simultaneous ascending auction was used.

Recent experience with auctions in other industries

Over the last ten years, there has been a great advance in the development of methods for auctioning many related items. Innovative auction designs have been proposed and applied to allocation problems in several industries. The auction of radio spectrum is one important example, but these methods have been adopted in several industries, such as energy and transportation.

Simultaneous ascending auction

The simultaneous ascending auction is one of the most successful methods for auctioning many related items. It was first introduced in U.S. spectrum auctions in July 1994, and later used in dozens of spectrum auctions worldwide, resulting in revenues in excess of $200 billion (Cramton 2002).

The simultaneous ascending auction is a natural generalization of the English auction when selling many items. The key features are that all the items are up for auction at the same time, each with a price associated with it, and the bidders can bid on any of the items. The bidding continues until no bidder is willing to raise the bid on any of the items. Then the auction ends with each bidder winning the items on which it has the high bid, and paying its bid for any items won.

The reason for the success of this simple procedure is the excellent "price discovery" it affords.[4] As the auction progresses bidders see the tentative price information and condition subsequent bids on this new information. Over the course of the auction, bidders are able to develop a sense of what the final prices are likely to be, and can adjust their purchases in response to this price information. To the extent price information is sufficiently good and the bidders retain sufficient flexibility to shift toward their best package, the exposure problem is mitigated -- bidders are able to piece together a desirable package of items, despite the constraint of bidding on individual items rather than packages. Moreover, the price information helps the bidders focus their valuation efforts in the relevant range of the price space.

Auctions have become the preferred method of assigning spectrum and most have been simultaneous ascending auctions. (See Cramton 1997 and Milgrom 2004 for a history of the auctions.) There is now substantial evidence that this auction design has been successful (Cramton 1997, McAfee and McMillan 1996), with revenues often exceeding industry and government estimates. The simultaneous ascending auction may be partially responsible for the large revenues. By revealing information in the auction process, bidder uncertainty is reduced, and the bidders safely can bid more aggressively. Also, revenues may increase to the extent the design enables bidders to piece together more efficient packages of items.

Despite the general success, the simultaneous ascending auctions have experienced a few problems from which one can draw important lessons (Cramton and Schwartz 2002). One basic problem is the simultaneous ascending auction's vulnerability to revenue-reducing strategies in situations where competition is weak. Bidders have an incentive to try to work together to reduce their demands in order to keep prices low, and to use bid signaling strategies to coordinate on a split of the items.

A second problem in the early U.S. auctions arose from overly generous installment payment terms for small businesses. This led to speculative bidding. Winning prices were well above subsequent market prices, and most firms defaulted on the installments and went into bankruptcy. The end result was that substantial portions of the mobile wireless capacity lay fallow for nearly ten years. Some 3G auctions in Europe (notably the U.K. and German auctions) also ended at prices well in excess of subsequent market prices. However, the European auctions did not allow installment payments, so the outcome was simply a wealth transfer from the shareholders of the telecommunications companies to the taxpayers.

Simultaneous clock auction

A variation of the simultaneous ascending auction is the simultaneous clock auction. The critical difference is that bidders simply respond with *quantities* desired at prices specified by the auctioneer. Clock auctions are especially effective in auctioning many divisible goods, like electricity, but the approach also works well for indivisible items like oil licenses. There is a price "clock" for each item indicating its tentative price. Bidders express the blocks desired at the current prices. For those blocks with excess demand the price is raised and bidders again express their desired blocks at the new prices. This process continues until supply just equals demand. The tentative prices and assignments then become final.

Clock auctions have been used with great success in many countries to auction electricity, gas, pollution allowances, and radio spectrum. Participants value the simplicity and price discovery of the auction. Furthermore, auction theory suggests that in a competitive setting, if bidding is continuous, then the clock auction is efficient with prices equal to the competitive equilibrium (Ausubel and Cramton 2004). This provides a strong rationale to adopt this auction format.

If bidding is not continuous (in the sense that there are discrete rounds) then issues of bid increments, ties, and rationing are important. But there is an easy solution to these problems in which bidders are allowed to express their demands for all points along the line segment between the start of round prices and the end of round prices. Allowing a rich expression of preferences within a round makes bid increments, ties, and rationing less important. Since preferences for intermediate prices can be expressed, the efficiency loss associated with the discrete increment is less, so the auctioneer can choose a larger bid increment, resulting in a faster and less costly auction process.

The problem of non-competitive settings is more difficult to handle, but there are some things that can be done. Although some auction settings approximate the ideal of perfect competition, most do not. In the U.S. oil auctions, especially in recent years when more marginal tracts have been offered, it is common for tracts to receive one or zero bids. In such a setting, tacit collusion is a real concern with the dynamic auction. One solution is to limit the amount of information that the bidders receive after each round of the auction. In doing this, the auctioneer can enhance the desirable properties of price and assignment discovery, while limiting the scope for collusive bidding. In the clock auction, this is done by only reporting the total quantity demanded for each block, rather than all the bids and bidder identities, as is commonly done in the simultaneous ascending auction.

Details matter

Not all auctions are successful. The most common source of failure is a lack of participation. Sometimes this is because what is being sold has little value. Other times the lack of competition is the result of a poor auction process, for example the product is ill-defined, the marketing is inadequate, or the political risks are too great. Recognition of the needs of the bidders is critical in getting participation. An important lesson is that careful planning and design are essential to maximizing results and that this can often be achieved through a prior process of consultation with prospective bidders. These efforts can translate into billions of dollars in higher revenues.

A practical package auction

We have seen that there are theoretical reasons to expect that some version of a simultaneous ascending auction will make sense for auctioning oil rights; we have seen, furthermore, that this type of auction design has been very successful in other industries. In this section, I now describe a practical method for auctioning many related items. The version I describe -- the clock-proxy auction (Ausubel, Cramton, and Milgrom 2006; hereafter ACM) -- also allows for package bids and so is especially appropriate to some oil and gas settings. I then turn to discuss variations for situations where packaging issues are less important. All methods are described with oil rights auctions in mind. The items sold are licenses to explore and develop specified geographic blocks. The bidder expresses quantities of either 0 ("no") or 1 ("yes") for each block offered.

The method combines two auction formats -- the clock auction and the proxy auction -- to produce a hybrid with the benefits of both.

The *clock auction*, as described earlier, is an iterative auction procedure in which the auctioneer announces prices, one for each of the items being sold. The bidders then indicate the licenses desired at the current prices. Prices for items with excess demand then increase, and the bidders again express quantities at the new prices. This process is repeated until there are no items with excess demand.

The *ascending proxy auction* is a particular package bidding procedure with desirable properties (Ausubel and Milgrom 2002). The bidders report values to their respective proxy agents. The proxies are in fact algorithms that receive information from the bidders just once and are then programmed to bid in the interests of the bidder. They iteratively submit package bids on behalf of the bidders, selecting the best profit opportunity for a bidder given the bidder's inputted values. For example if the proxy is told by a bidder that a block is worth $100M, the proxy will bid up to $100M in order to get the block but will never bid more than is needed to secure the block. The auctioneer then selects the provisionally winning bids that maximize revenues. This process continues until the proxy agents have no new bids to submit.

The clock-proxy auction begins with a clock phase and ends with a final proxy round. All bids are kept live throughout the auction. There are no bid withdrawals and the bids of a particular bidder are mutually exclusive. Finally, there is an activity rule throughout the clock phase and between the clock phase and the proxy round.

Clock phase

The clock phase has several important benefits. First, it is simple for the bidders. At each round, the bidder simply expresses the set of licenses desired at the current prices. Additive pricing means that it is trivial to evaluate the cost of any package -- it is just the sum of the prices for the selected licenses. Limiting the bidders' information to a reporting of the excess demand for each item removes much strategizing. Complex bid signaling and collusive strategies are eliminated, as the bidders cannot see individual bids, but only aggregate information. Second, the clock phase produces highly useable price discovery, because of the item prices. With each bidding round, the bidders get a better understanding of the likely prices for relevant packages. This is essential information in guiding the bidders' decision making. Bidders are able to focus their valuation efforts on the most relevant portion of the price space. As a result, the valuation efforts are more productive. Bidder participation costs fall and efficiency improves.

The weakness of the clock auction is its use of additive pricing at the end of the auction. This means that, to the extent that there is market power, bidders will have an incentive to engage in demand reduction to favorably impact prices. This demand reduction implies that the auction outcome will not be fully efficient (Ausubel and Cramton 2002). The proxy phase will however eliminate this inefficiency.

There are several design choices that will improve the performance of the clock phase, when packaging issues are important. Good choices can avoid the exposure problem, improve price discovery, and handle discrete rounds.

Avoiding the exposure problem

To avoid the exposure problem, the clock phase can allow a bidder to drop a license so long as the price has increased on some license the bidder was demanding. This flexibility is needed in the case of complements. The bidder may want to drop a license when the price of a complementary license increases.

With this rule, the clock auction becomes a package auction. For each price vector, the bidder expresses the package of licenses desired without committing itself to demanding any smaller package.

The disadvantage of this rule is that the clock phase may end with a substantial number of unsold licenses. However, this undersell will be resolved in the proxy phase.

Improving price discovery

In auctions with more than a few items, the sheer number of packages that a bidder might buy makes it impossible for bidders to determine all their values in advance. Bidders adapt to this problem by focusing most of their attention on the packages that are likely to be valuable relative to their forecast prices. A common heuristic to forecast package prices is to estimate the prices of individual items and to sum these up over the elements in the package. Clock auctions with individual prices assist bidders in this *price discovery* process.

Price discovery is undermined to the extent that bidders misrepresent their demands early in the auction. One possibility is that bidders will choose to underbid in the clock phase, hiding as a "snake in the grass" to conceal their true interests from their opponents. To limit this form of insincere bidding, the U.S. Federal Communications Commission (FCC) introduced the Milgrom-Wilson activity rule, and similar activity rules have since become standard in both clock auctions and simultaneous ascending auctions. In its most typical form, a bidder desiring large quantities at the end of the auction must bid for quantities at least as large early in the auction, when prices are lower.

The most common activity rule in clock auctions is monotonicity in quantity. As prices rise, quantities cannot increase. This means that bidders must bid in a way that is consistent with a weakly downward sloping demand curve. This works well when auctioning a single product, but is overly restrictive when there are many different products. If the products are substitutes, it is natural for a bidder to want to shift quantity from one product to another as prices change, effectively arbitraging the price differences between substitute products.
A weaker activity requirement is a monotonicity of a bidder's *aggregate* quantity. This allows flexibility in shifting among licenses. This is the basis for the FCC's activity rule. A weakness of this rule is that it assumes that quantities are readily comparable. Oil licenses, however, are not comparable. For example, the area of the block is a poor measure of quantity.

ACM propose an alternative activity rule based on revealed preference that does not require any aggregate quantity measure. This rule is suitable for bonus bidding, but not for bidding on production shares. The rule is derived from standard consumer theory and works as follows. Consider any two times, denoted s and t ($s < t$). Let p^s and p^t be the price vectors at these times. Let x^s and x^t be the associated demands of some bidder at these two points in time, and assume that x^s is different from x^t. Finally let $v(x)$ be the bidder's value of a given package x. Unfortunately $v(x)$ is not known, nonetheless we can derive a very simple condition on the relationship between the prices and the quantities if x^s and x^t reflect the demands of a sincere bidder, even though v is not known. With a price vector p^s and a vector of demands for blocks x^s (x^s is in this case a set of 0s and 1s), the amount that a buyer pays is given by $p^s \cdot x^s$. Similarly the total cost of a vector x^t at prices p^s is given by $p^s \cdot x^t$. Now, a sincere bidder prefers x^s to x^t when prices are p^s and prefers x^t to x^s when prices are p^t. This means that:

$$v(x^s) - p^s \cdot x^s \geq v(x^t) - p^s \cdot x^t$$

and

$$v(x^t) - p^t \cdot x^t \geq v(x^s) - p^t \cdot x^s.$$

Adding these two inequalities together yields the *revealed preference activity rule*:

$$(p^t - p^s) \cdot (x^t - x^s) \leq 0.$$

This inequality gives a condition on the prices and quantities that should hold for a sincere bidder without using any information about v. This condition can be used as a rule in auction design: at every time t, the bidder's demand x^t must satisfy this condition for all times $t>s$. Straightforward bidding -- bidding on the most profitable package at every instant -- will always satisfy the condition.

Handling discrete rounds

As described above, discrete bidding rounds are handled with intra-round bidding, enabling the bidder to express quantity reductions at any prices between the start-of-round price and the end-of-round price. This allows the use of much larger bid increments without much loss in efficiency. In this way, the auctioneer can better control the pace of the auction, which is important here given the large uncertainty in block values.

Proxy phase

Like the clock auction, the proxy auction is based on package bids. However, the incentives are quite different. The main difference is the absence of additive prices on individual items. Only packages are priced -- and the prices may be bidder specific. This weakens price discovery, but the proxy phase is not about price discovery. It is about providing the incentives for efficient assignment. All the price discovery occurs in the clock phase. The second main difference is that the bidders do not bid directly in the proxy phase. Rather, they submit values to the proxy agents, who then bid on their behalf using a specific bidding rule. The proxy agents bid straightforwardly to maximize profits. The proxy phase is a last-and-final opportunity to bid.

The proxy auction works as follows (see Ausubel and Milgrom 2002). Each bidder reports his values to a proxy agent for all packages that the bidder is interested in. Budget constraints can also be reported. The proxy agent then bids in an ascending package auction on behalf of the real bidder, iteratively submitting the allowable bid that, if accepted, would maximize the real bidder's profit (value minus price), based on the reported values. The auction in theory is conducted with infinitesimally small bid increments. After each round, provisionally winning bids are determined that maximize seller revenue from compatible bids. All of a bidder's bids are kept live throughout the auction and are treated as mutually exclusive. The auction ends after a round with no new bids. For more information on the practicalities of how to implement the proxy phase, see the description in Day and Raghavan (2004).

The advantage of this format is that it produces an outcome for which there is no *other* allocation that the seller and any collection of bidders all prefer (where "preference" is given in terms of the reported preferences of the bidders). In economic theory, we say that such an outcome is in the "core." In principle there can be many outcomes in the core, but every outcome that is in the core has the property that the seller would not be able to find an alternative group of buyers that would be willing, collectively to pay more for the goods on sale (see Ausubel and Milgrom 2002, Parkes and Ungar 2000). Hence under this method, the seller earns competitive revenues. The fact that the outcome is in the core also implies that it is *efficient*. In particular this means that the proxy auction is not subject to the inefficiency of demand reduction: no bidder can ever reduce the price it pays for the package it wins by withholding some of its losing bids for other packages.

The clock-proxy auction

The clock-proxy auction is a recent innovation in auction design. Although it has been used successfully in a spectrum auction in a developing country (Trinidad and Tobago, June 2005), it has not yet been applied to oil rights auctions. The clock-proxy auction begins with a clock auction for price discovery and concludes with the proxy auction to promote efficiency.

The clock auction is conducted with the revealed-preference activity rule until there is no excess demand on any item. The market-clearing item prices determine the initial minimum bids for all packages for all bidders. Bidders then submit values to proxy agents, who bid to maximize profits, subject to a relaxed revealed-preference activity rule (see ACM for details). The relaxed revealed-preference activity rule operates also between the rounds and so bidders that failed to bid aggressively in the clock stage are constrained on how aggressively they can behave in the proxy stage. The bids from the clock phase are kept live as package bids in the proxy phase. All of a bidder's bids, both clock and proxy, are treated as mutually exclusive. Thus, the auctioneer obtains the provisional winning bids after each round of the proxy phase by including all bids -- those submitted in the clock phase as well as those submitted in the proxy phase -- in the winner determination problem and by selecting at most one provisional winning bid from every bidder. As usual, the proxy phase ends after a round with no new bids.

Why include the clock phase?

The clock phase provides price discovery that bidders can use to guide their calculations in the complex package auction. At each round, bidders are faced with the simple and familiar problem of expressing demands at specified prices. Moreover, because there is no exposure problem, bidders can bid for synergistic gains without fear. Prices then adjust in response to excess demand. As the bidding continues, bidders get a better understanding of what they may win and where their best opportunities lie.

The case for the clock phase relies on the idea that it is costly for bidders to determine their preferences. The clock phase, by providing tentative price information, helps focus a bidder's decision problem. Rather than consider all possibilities from the outset, the bidder can instead focus on cases that are important given the tentative price and assignment information. Although the idea that bidders can make information processing decisions in auctions is valid even in auctions for a single good (Compte and Jehiel 2002), its importance is magnified when there are many goods for sale, because the bidder's decision problem is then much more complicated. Rather than simply decide whether to buy at a given price, the bidder must decide which goods to buy and how many of each. The number of possibilities grows exponentially with the number of goods. Price discovery can play an extremely valuable role in guiding the bidder through the valuation process.

Price discovery in the clock phase makes bidding in the proxy phase vastly simpler. Without the clock phase, bidders would be forced either to determine values for all possible packages or to make uninformed guesses about which packages were likely to be most attractive. My experience with dozens of bidders suggests that the second outcome is much more likely; determining the values of exponentially many packages becomes quickly impractical with even a modest number of items for sale. Using the clock phase to make informed guesses about prices, bidders can focus their decision making on the most relevant packages. The bidders see that they do not need to consider the vast majority of options, because the options are excluded by the prices established in the clock phase. The bidders also get a sense of what packages are most promising, and how their demands fit in the aggregate with those of the other bidders.

In competitive auctions where the items are substitutes and competition is strong, we can expect the clock phase to do most of the work in establishing prices and assignments -- the proxy phase would play a limited role. When competition is weak, demand reduction may lead the clock phase to end prematurely, but this problem is corrected at the proxy stage, which eliminates incentives for demand reduction. If the clock auction gives the bidders a good idea of likely package prices, then expressing a simple approximate valuation to the proxy is made easier.

Why include the proxy phase?

The main advantage of the proxy phase is that it pushes the outcome toward the core, that is, toward an efficient allocation with competitive payoffs for the bidders and competitive revenues for the seller.

In the proxy phase, there are no incentives for demand reduction. A large bidder can bid for large quantities without the fear that doing so will adversely impact the price the bidder pays.

The proxy phase also mitigates collusion. Any collusive split of the items established in the clock phase can be undone in the proxy phase. The relaxed activity rule means that the bidders can expand demands in the proxy phase. The allocation is still up for grabs in the proxy phase.

A natural concern with the proxy phase is that it may discourage bidding in the clock phase. The activity rule that operates between the two phases is essential in mitigating this possibility. Bidders bid aggressively in the clock phase, knowing that a failure to do so will limit their options in the proxy phase.

Implementation issues

We briefly discuss three important implementation issues.

Confidentiality of values

One practical issue with the proxy phase is confidentiality of values. Bidders may be hesitant to bid true values in the proxy phase, fearing that the auctioneer would somehow manipulate the prices with a "seller shill"[5] to push prices all the way to the bidders' reported values. Steps need to be taken to assure that this cannot happen. A highly transparent auction process helps to assure that the auction rules are followed. In fact, there is no reason that the auctioneer needs to be given access to the high values. Only the computer need know. Auction software can be tested and certified to be consistent with the auction rules. At the end of the auction, the auctioneer can report all the bids. The bidders can then confirm that the outcome was consistent with the rules.

Price increments in the clock phase

When auctioning many items, one must take care in defining the price adjustment process. This is especially true when some goods are complements. Intuitively, the clock phase performs best when each item clears at roughly the same time. Thus, the goal should be to come up with a price adjustment process that reflects relative values as well as excess demand. Moreover, the price adjustment process effectively is resolving the "threshold problem" by specifying who should contribute what as the clock ticks higher.[6] To the extent that prices adjust with relative values, the resolution of the threshold problem will be more successful.

One simple approach is to build the relative value information into the initial starting prices. Then use a *percentage* increase rather than a fixed increment increase, based on the extent of excess demand. For example, the percentage increment could vary linearly with the excess demand, subject to a lower and upper limit. This will make it more likely that the demand for the different items clears at about the same time.

Expression of proxy values

Even with the benefit of the price discovery in the clock phase, expressing a valuation function in the proxy phase may be difficult for bidders. When many items are being sold, the bidder will need a tool to facilitate translating preferences into proxy values. The best tool will depend on the circumstances. The seller can lower bidder participation costs by providing a useful tool to express valuations.

At a minimum, the tool will allow an additive valuation function. The bidder submits its maximum willingness to pay for each license. The value of a package is then found by adding up the values on each item in the package. This additive model ignores all value interdependencies across items; it assumes that the value for one item is independent of what other items are won. Although globally (across a wide range of packages) this might be a bad assumption, locally (across a narrow range of packages) this might be a reasonable approximation, especially in the setting of oil rights. Hence, provided the clock phase has taken us close to the equilibrium, so the proxy phase is only doing some fine-tuning of the clock outcome, then such a simplistic tool may perform reasonably well. And of course it performs very well when bidders actually have additive values.

The bidders' business plans are a useful guide to determine how best to structure the valuation tool in a particular setting. Business plans are an expression of value to investors. Although the details of the business plans are not available to the auctioneer, one can construct a useful valuation tool from understanding the basic structure of these business plans.

Alternative auction formats and recommendations

It is not possible to specify one "best" design -- the best approach depends on the setting. The clock-proxy auction as described above is an excellent choice in settings where packaging issues are important. In other settings, variations are worth considering. The variations depend on how four features are handled (see Table 5.1).

Table 5.1

Designing a clock-proxy auction

Feature:	1. Clock bidding	2. Activity rule	3. Proxy bidding	4. Pricing in proxy phase
Options:	a. Package bids b. Individual license bids c. None	a. Revealed preference b. License-by-license monotonicity	a. Package bids b. Individual license bids c. None	a. Bidder-optimal core b. Pay-as-bid

Note: This table lists the main options available to auction designers for four key components of the clock-proxy auction.

The standard clock-proxy auction is defined by the first option (a) for each issue in Table 5.1: clock bidding for packages with the revealed preference activity rule, followed by a final proxy round with package bids and bidder-optimal core pricing. This is a sensible choice when packaging issues, value interdependencies and price discovery are important aspects of the setting. This approach is the most difficult to implement, but accommodates the richest set of bidder valuations.

At the other extreme is the U.S. offshore approach, which is simultaneous seal-bid for individual licenses with pay-as-bid pricing (1c, 3b, 4b). This approach makes sense if there are no packaging issues (for example, additive values), little value interdependencies, weak competition, and potentially large asymmetries among the bidders. Although this method is easy to implement, it is problematic for bidders unless values really are additive.

Another variation, close to the U.S. approach, has clock bidding on individual licenses, a license-by-license activity rule, and no proxy bidding (1b, 2b, 3c). This effectively is a simultaneous ascending auction version of the U.S. approach. This is sensible in settings where packaging is of only minor importance (nearly additive values), but value interdependencies make price discovery important. This approach also works best when competition is not too weak and bidder asymmetries are not too large.

A similar variation, close to the U.S. spectrum auctions is clock bidding on individual licenses, a revealed preference activity rule, and no proxy round (1b, 2a, 3c). This would work well when there are moderate packaging issues and value interdependencies. The approach has good price discovery and does allow bidders to piece together desirable packages of licenses. The format improves on the U.S. spectrum auctions in two respects. Tacit collusion is mitigated with the use of clocks and only reporting excess demand, rather than all bids. Efficient packaging is facilitated with the revealed preference activity rule. This method is easy to implement and yet accommodates a richer set of valuations.

A final variation, related to the Anglo-Dutch format[7] (Klemperer 2002), has clock bidding on individual licenses, a revealed preference activity rule, and a proxy round with individual license bids and pay-as-bid pricing (1b, 2a, 3b, 4b). However, in this variation, the price clock stops when demand falls to two on the license, so there is still excess demand. The excess demand is then resolved in the simultaneous pay-as-bid proxy round. This approach is well-suited to situations where packaging is of minor importance (nearly additive values), but value interdependencies make price discovery valuable, and competition is weak with potentially large bidder asymmetries. The approach enjoys some of the price discovery benefits of the dynamic methods, but handles weak competition and bidder asymmetries better than the approach without a last-and-final round.

The approaches are summarized in Table 5.2.

Table 5.2

Alternative Auction Approaches

Interdependence of Valuations Across Bidders	Interdependence of Valuations Across Blocks: Additive values	Nearly additive values	Substitutes and mild complements	Complex structure of substitutes and complements
Private values	**First-Price Sealed-Bid** Simultaneous sealed-bid Pay-as-bid pricing **Note**: Easiest to implement No price discovery Handles weak competition Handles bidder asymmetries.	**Anglo-Dutch Clock** (see below	**Clock with Switching** (see below)	**Clock-Proxy** (see below)
Mostly private values	**Anglo-Dutch Clock** (see right)	**Anglo-Dutch Clock** Clock individual bids (stops with demand = 2) Revealed preference activity rule; Proxy with individual bids Pay-as-bid pricing **Note**: Harder to implement Some price discovery Handles weak competition Handles bidder asymmetries	**Clock with Switching** (see below)	**Clock-Proxy** (see below)
Interdependent values	**Clock No Switching** (see right)	**Clock No Switching** Clock individual bids License-by-license activity rule **Note**: Easy to implement Good price discovery with nearly additive values Handles production shares	**Clock with Switching** Clock individual bids Revealed preference activity rule No final proxy round. **Note**: Harder to implement Very good price discovery	**Clock-Proxy** Clock package bids Revealed preference activity rule; Proxy package bids Bidder-optimal core pricing. **Note**: Hardest to implement Excellent price discovery Excellent efficiency Competitive revenues

For settings where there are sets of licenses with substantially different value structures, it makes sense to use different formats with different sets of licenses. For example, a country may have 12 wildcat blocks that are excellent prospects, 36 drainage blocks that are good to excellent prospects, and 200 blocks that are marginal prospects. The excellent prospects could be done as a standard clock-proxy, the drainage blocks as an Anglo-Dutch, and the marginal prospects as a first-price sealed-bid. With this approach the clock-proxy auction is not complicated by the great number of drainage and marginal blocks. Moreover, the drainage blocks may have large asymmetries among the bidders as a result of private drilling information from neighboring blocks. The Anglo-Dutch design handles these asymmetries well. Finally, additive values are probably a good assumption on marginal prospects and in any event the economic loss from the less efficient first-price sealed-bid approach is not great when auctioning marginal licenses. Alternatively, since implementing three different formats is probably too much, the country could split the blocks into two sets: those with high prospects and those with low prospects. The first-price sealed-bid format could be used for the low-prospect blocks and one of the dynamic formats could be used for the high-prospect blocks.

Libya and Venezuela reconsidered

Although the 2005 Libya auction and 1996 Venezuela auction were successful, they could probably be improved. The Libya auction, using simultaneous sealed-bids, prevented both price discovery and efficient packaging. The Venezuela auction, using sequential sealed-bids, allowed only minimal price discovery and packaging. In both auctions, competition was anticipated to be strong. Values included both private and common elements, although the common elements were more important. Values probably were nearly additive, although bidders likely faced budget and risk constraints given the size of the commitment.

In such a setting, a simultaneous clock auction is desirable, and especially simple given the small number of blocks. Bids would be over the production share. In the case of Venezuela, the 50 percent cap on production share could be dropped and the terms could be adjusted so that the government has a share in the development capital expense, thereby improving the development incentives without limiting the production share. A license-by-license activity rule (no switching) is desirable given the bidding is on production shares. Under this rule, once a bidder stops bidding on a block, the bidder cannot return to the block at higher production shares. This simple rule allows price discovery and some degree of packaging.

Conclusion

Auctions are a desirable method of assigning and pricing scarce oil rights. A well-designed auction encourages participation through a transparent competitive process. The design promotes both an efficient assignment of the rights and competitive revenues for the seller.

I find that a variety of auction formats are suitable for auctioning oil rights. The best auction format depends on the particular setting, especially the structure of bidder preferences and the degree of competition. When bidders have additive values and competition is weak, a simultaneous first-price sealed-bid auction may be best, especially if the blocks are marginal prospects (relatively low value). When bidders have nearly additive values and competition is stronger, then one of the clock auctions should be considered. This approach will improve price discovery and reduce bidder uncertainty, improving efficiency and revenues. Finally, for high-value blocks in which packaging issues are important (bidders care about the particular package of blocks won), a clock-proxy auction is appropriate. The clock-proxy auction has excellent price discovery and handles complex bidder preferences involving substitutes and complements. The clock-proxy auction does well on both efficiency and revenue grounds.

In closing, I emphasize a theme raised elsewhere in this volume (see especially Chapters 2 and 3): Regardless of the auction format, a critical element of the design is defining *what* is being bid. The *what* element is just as important as the *how* element. Possibilities include bonus bids, royalty rates, or production shares. Bidding on production shares, rather than bonuses, for example, typically increases Government Take by reducing oil company risk and fears of expropriation. More generally, these contract terms determine key features such as the allocation of risk between country and oil company, the cash flows over time, and the incentives for exploration and development.

References

Ausubel, Lawrence M. and Peter Cramton. 2004. "Auctioning Many Divisible Goods." *Journal of the European Economic Association* 2: 480-493.

Ausubel, Lawrence M., Peter Cramton, and Paul Milgrom. 2006. "The Clock-Proxy Auction: A Practical Combinatorial Auction Design." In *Combinatorial Auctions*, Peter Cramton, Yoav Shoham, and Richard Steinberg, eds., 115-138 Cambridge: MIT Press.

Ausubel, Lawrence M. and Paul Milgrom. 2002. "Ascending Auctions with Package Bidding." *Frontiers of Theoretical Economics* 1: 1–45. www.bepress.com/bejte/frontiers/vol1/iss1/art1.

Compte, Olivier and Philippe Jehiel. 2002. "Auctions and Information Acquisition: Sealed-bid or Dynamic Formats?" Working Paper, CERAS-ENPC.

Cramton, Peter. 1997. "The FCC Spectrum Auctions: An Early Assessment." *Journal of Economics and Management Strategy*. 6: 431-495.

Cramton, Peter. 2002. "Spectrum Auctions." In *Handbook of Telecommunications Economics*, Martin Cave, Sumit Majumdar, and Ingo Vogelsang, eds., pp. 605-639. Amsterdam: Elsevier Science B.V.

Cramton, Peter and Jesse Schwartz. 2002. "Collusive Bidding in the FCC Spectrum Auctions." *Contributions to Economic Analysis & Policy* 1: 1. www.bepress.com/bejeap/contributions/vol1/iss1/art11.

Cramton, Peter, Yoav Shoham, and Richard Steinberg. 2006. *Combinatorial Auctions*, MIT Press.

Day, Robert W. and S. Raghavan. 2004. "Generation and Selection of Core Outcomes in Sealed-Bid Combinatorial Auctions." Working Paper, University of Maryland.

Hendricks, Kenneth and Robert H. Porter. 1996. "The Timing and Incidence of Exploratory Drilling on Offshore Wildcat Tracts." *American Economic Review* 86: 388-407.

Hendricks, Kenneth, Robert H. Porter, and Bryan Boudreau. 1987. "Information, Returns, and Bidding Behavior in OCS Auctions: 1954-1969." *Journal of Industrial Economics* 35: 517-542.

Johnston, Daniel. 2005. "Tough Terms -- No Surprises: Libya EPSA IV License Round - 29 January 2005." White Paper, Daniel Johnston & Co. Inc.

Johnston, Daniel and David Johnston. 2001. "Kashagan and Tengiz - Castor and Plux." White Paper, Daniel Johnston & Co. Inc.

Klemperer, Paul. 2002. "What Really Matters in Auction Design." *Journal of Economic Perspectives* 16: 169-189.

Maskin, Eric and John Riley. 2000. "Asymmetric Auctions." *Review of Economic Studies* 67: 439-454.

McAfee, R. Preston and John McMillan. 1996. "Analyzing the Airwaves Auction." *Journal of Economic Perspectives* 10: 159-176.

Milgrom, Paul. 1987. "Auction Theory." In *Advances in Economic Theory - Fifth World Congress*, Truman Bewley, ed., 1-31 Cambridge: Cambridge University Press.

Milgrom, Paul. 2004. *Putting Auction Theory to Work*. Cambridge: Cambridge University Press.

Milgrom, Paul and Robert J. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50: 1089-1122.

Parkes, David C. and Lyle H. Ungar. 2000. "Iterative Combinatorial Auctions: Theory and Practice." Proceedings of the 17th National Conference on Artificial Intelligence (AAAI-00), 74–81.

Porter, Robert H. 1995. "The Role of Information in U.S. Offshore Oil and Gas Lease Auctions." *Econometrica* 63: 1-28.

Sunley, Emil M., Thomas Baunsgaard, and Dominique Simard. 2002. "Revenue from the Oil and Gas Sector: Issues and Country Experience." IMF Conference Paper.

---

[*] Economics Department, University of Maryland, College Park, MD 20742, USA. Correspondence to cramton@umd.edu. I gratefully acknowledge the support of National Science Foundation Grants SES-01-12906, IIS-02-05489, and SES-05-31254.

[1] Sealed-bid auctions are static auctions in which the bidders simultaneously submit bids in sealed envelopes to the auctioneer. The auctioneer then orders the bids from highest to lowest. The highest bidder wins the item. In a first-price auction, the highest bidder pays its bid. In a second-price auction, the highest bidder pays the

second-highest bid for the item. The English ascending and Dutch descending are dynamic auctions that occur in a series of rounds or with a continuous price 'clock.' In the English auction, the bidders have the opportunity to improve their bids until no bidder is willing to bid any higher; the high bidder then pays its bid. In the Dutch auction, the auctioneer begins with a high price that no bidder is willing to accept; the auctioneer then reduces the price until a bidder indicates acceptance; this bidder wins the item and pays the last price called out by the auctioneer.

[2] Although, a likely exception to this recommendation is for drainage blocks in which one bidder has much better information about value.

[3] Government Take is a calculated estimate of the government's share in the project profits, given all the contract terms. Production share is simply the split of the oil revenue between government and the oil company. Thus, it is perfectly reasonable for Government Take to be 92 percent even though the production share is 50 percent.

[4] Price discovery is a feature of dynamic auctions in which tentative price information is reported to bidders, giving bidders the opportunity to adjust subsequent bids based on the price information.

[5] A seller shill is a fake bidder created by the seller. The shill bids in a way to increase auction revenues. Shills are especially a problem in ascending auctions and second-price auctions.

[6] The threshold problem is the problem that allowing package bids may favor bidders seeking larger packages, because small bidders do not have the incentive or capability to unilaterally top the tentative winning bids of a large bidder.

[7] The Anglo-Dutch format is an ascending clock auction that stops while there is still excess demand, at which point a final first-price sealed-bid round is conducted among those bidders that still active at the end of the clock phase.