Hurwicz, Leonid (1973). "The Design of Mechanisms for Resource Allocation." *American Economic Review* **63**(2): 1–30.

Milgrom, Paul and Nancy Stokey (1982). "Information, Trade and Common Knowledge." *Journal of Economic Theory* **26**: 17–27.

Myerson, Roger B. (1981). "Optimal Auction Design." *Mathematics of Operations Research* **6**(1): 58–73.

Ortega-Reichert, Armando (1968). *Models for Competitive Bidding under Uncertainty.* Stanford, CA: Stanford University.

Wilson, Robert (1969). "Competitive Bidding with Disparate Information." *Management Science* **15**(7): 446–448.

CHAPTER TWO

Vickrey–Clarke–Groves Mechanisms

This chapter describes the important contributions of Vickrey, Clarke, and Groves (VCG) to the theory of mechanism design. Vickrey (1961) analyzed a situation in which bidders compete to buy or sell a collection of goods. Later, Clarke (1971) and Groves (1973) studied the public choice problem, in which agents decide whether to undertake a public project – e.g. construction of a bridge or highway – whose cost must be borne by the agents. This latter analysis formally includes any choice from a finite set. In particular, it includes the Vickrey analysis for the case of discrete assets. We limit attention in this chapter to the case of finite choice sets, to bypass technical issues associated with infinite choice sets, particularly issues associated with the existence of a best choice.

The VCG analysis has become an important standard. It is the work by which nearly all other mechanism design work is judged and in terms of which its contribution is assessed. As we will see in later chapters, there are deep and surprising connections between the VCG theory and many other parts of auction theory.

2.1 Formulation

We begin the theoretical development in this section by introducing notation and defining direct mechanisms and VCG mechanisms.

Thus, let $N = \{0, \ldots, n\}$ denote the set of participants, with participant 0 being the mechanism operator. Let X denote the set of possible decisions with typical element x. For chapters 2–5, we assume that the set of participants is exogenously given and omit any analysis of the incentives to participate. An outcome is a pair (x, p) describing a decision x and a vector of positive or negative payments $p = (p^0, p^1, \ldots, p^n)$ by the participants. For example, in a first-price sealed-bid auction, the decision x is a vector where $x^i = 1$ if agent i gets the object and 0 otherwise. The associated vector of payments is p, where $p^i = b^i = -p^0$ if i bids b^i and wins, and in that case, $p^j = 0$ for the other bidders.

For most of our analysis, we also assume that each participant i values outcomes according to $u^i((x, p), \vec{t}) \equiv v^i(x, t^i) - p^i$, that is, i's payoff corresponding to outcome (x, p) is i's value of the decision x, which depends only on i's own type t^i, minus the payment that i must make. This *quasi-linear* specification of the utility function plays an indispensable role in the formal analysis of this chapter. The assumption of quasi-linearity implies that bidders are able to make any cash transfers described by the mechanism, that there exists a cash transfer that exactly compensates any individual for any possible change in outcomes, and that redistributing wealth among the participants would not change this compensatory transfer. These assumptions represent better modeling approximations for some situations than for others. For example, if the bidders are firms with ample liquidity, the assumptions might be a very good approximation of reality, but if they are consumers with significant credit constraints that apply to the transactions, then the assumptions might be an unacceptably bad fit.

Recall that "performance" means the function that maps environments into outcomes. Given our assumptions about the two-part description of outcomes, the performance of any mechanism can be also described in two parts. The *decision performance* maps types into decisions x, whereas the *transfer performance* maps types into payments or transfers. When the decision x allocates goods, we sometimes also call x the *allocation performance*.

The VCG analysis sometimes attempts to achieve efficient performance subject to the constraint that transfers add up to zero. Given the assumptions described above, a decision x is efficient if it maximizes the total value $\sum_{i \in N} v^i(x, t^i)$. For example, in an auction of a single good, the final allocation is efficient if it awards the good to the bidder who values it the most. In the models studied here, by construction, net payments always add up to zero, because the seller (or mechanism designer) receives any sums that the buyers (or bidders) pay.

In some publicly run auctions, the design objective is efficiency as defined above, although revenues (the total transfer to the mechanism designer) may also be an important goal. In private-sector auctions, revenues are always an important goal and often the only one.

Sometimes, the designer wants to run an auction in which $p^0 \equiv 0$, that is, in which there is never any net transfer to the auction designer. These *balanced-budget mechanisms* are useful, for example, in regulatory contexts where the regulator is not authorized to contribute or collect money from the regulated parties. They also arise in the theory of the firm, where the mechanism operator is similarly restricted. As we will see later, there is often a tension in mechanism design between achieving efficient outcomes and ensuring a balanced budget.

The VCG mechanisms are *incentive-compatible direct mechanisms.* This means that (1) $S = \Theta$ and that (2) the strategy profile $(\sigma^i(t^i) = t^i)_{i \in N}$ is an equilibrium. In words, the first condition means that each participant is required to report a possible type to the mechanism operator. We will sometimes speak of direct mechanisms as being pairs (x, p), leaving the strategy set implicit. The second condition, *incentive compatibility,* means that reporting one's type truthfully is an equilibrium according to whatever solution concept we have chosen. For VCG mechanisms, we focus on *dominant-strategy implementation,* so the relevant solution concept is that each participant plays a dominant strategy.

One appeal of incentive-compatible direct mechanisms is that they spare participants the need for elaborate strategic calculations: truthful reporting serves each participant's individual interest. Choosing dominant strategies as the solution concept, an incentive-compatible direct mechanism is one for which truthful reporting leads to as high a payoff as any other strategy for all possible types of opponents and all possible actions that these opponents may take. For example, as discussed in chapter 1, it is always optimal for a bidder in a second-price sealed-bid auction for a single good to bid his valuation. Moreover, this *truthful* bidding strategy is the only strategy that is always optimal, so it is a *dominant strategy.* Thus, the second-price auction is a dominant-strategy incentive-compatible direct mechanism.

The operator of a VCG mechanism uses the reported types to compute the maximum total value $V(X, N, \vec{t})$ and a corresponding total-value-maximizing decision $\hat{x}(X, N, \vec{t})$ as follows:

$$V(X, N, \vec{t}) = \max_{x \in X} \sum_{j \in N} v^j(x, t^j), \tag{2.1}$$

$$\hat{x}(X, N, \vec{t}) \in \arg\max_{x \in X} \sum_{j \in N} v^j(x, t^j). \tag{2.2}$$

One might think that such a direct approach would be doomed to failure, because each participant seems to have an incentive to misrepresent his preferences to influence the decision in his favor. However, the participants' incentives depend not only on the decision but also on the cash transfer, which is the clever and surprising part of the VCG mechanism.

The VCG mechanism eliminates incentives for misreporting by imposing on each participant the cost of any distortion he causes. The VCG payment for participant i is set so that i's report cannot affect the total payoff to the set of *other* parties, $N-i$. Notice that $0 \in N-i$, that is, the set includes the mechanism designer whose payoff is the mechanism's net receipts.

With this principle in mind, let us derive a formula for the VCG payments. To capture the effect of i's report on the outcome, we introduce a hypothetical *null report,* which corresponds to bidder i reporting that he is indifferent among the possible decisions and cares only about transfers. When i makes the null report, the VCG mechanism optimally chooses the decision $\hat{x}(X, N-i, t^{-i})$. The resulting total value of the decision for the set of participants $N-i$ would be $V(X, N-i, t^{-i})$, and the mechanism designer might also collect a payment $h^i(t^{-i})$ from participant i. Thus, if i makes a null report, the total payoff to the participants in set $N-i$ is

$$V(X, N-i, t^{-i}) + h^i(t^{-i}).$$

The VCG mechanism is constructed so that this same amount is the total payoff to those participants regardless of i's report. Thus, suppose that when the reported type profile is $\vec{t}$, i's payment is $\hat{p}^i(X, N, \vec{t}) + h^i(t^{-i})$, so that $\hat{p}^i(X, N, \vec{t})$ is i's additional payment over what i would pay if he made the null report. The decision $\hat{x}(X, N, \vec{t})$ generally depends on i's report, and the total payoff to members of $N-i$ is then $\sum_{j \in N-i} v^j(\hat{x}(X, N, \vec{t}), t^j) + \hat{p}^i(X, N, \vec{t}) + h^i(t^{-i})$. We equate this total value with the corresponding total value when i makes the null report:

$$\hat{p}^i(X, N, \vec{t}) + h^i(t^{-i}) + \sum_{j \in N-i} v^j(\hat{x}(X, N, \vec{t}), t^j)$$
$$= h^i(t^{-i}) + V(X, N-i, t^{-i}). \tag{2.3}$$

Using (2.1), we solve for the extra payment as follows:

$$\begin{aligned} \hat{p}^i(X, N, \vec{t}) &= V(X, N-i, t^{-i}) - \sum_{j \in N-i} v^j(\hat{x}(X, N, \vec{t}), t^j) \\ &= \sum_{j \in N-i} v^j(\hat{x}(X, N-i, t^{-i}), t^j) \\ &\quad - \sum_{j \in N-i} v^j(\hat{x}(X, N, \vec{t}), t^j). \end{aligned} \tag{2.4}$$

According to (2.4), if participant i's report leads to a change in the decision $\hat{x}$, then i's extra payment $\hat{p}^i(X, N, \vec{t})$ is specified to compensate the members of $N-i$ for the total losses they suffer on that account.

We now introduce some definitions:

Definition

1. A *Vickrey–Clarke–Groves (VCG) mechanism* $(\Theta, (\hat{x}, \hat{p}+h))$ is a direct mechanism in which $\hat{x}$ satisfies (2.2), $\hat{p}$ satisfies (2.4) (for all N, X, $\vec{t}$, and $i \in N$), and payments are determined by $\hat{p}^i(X, N, \vec{t}) + h^i(t^{-i})$.
2. A participant is *pivotal* if $\hat{x}(X, N, \vec{t}) \neq \hat{x}(X, N-i, t^{-i})$.
3. The *pivot mechanism* is the VCG mechanism in which $h^i \equiv 0$ for all $i \in N$.

In words, a participant is pivotal if consideration of his report changes the decision, compared to excluding the participant or attributing the null report to him. According to (2.4), if participant i is not pivotal, then $\hat{p}^i(X, N, \vec{t}) = 0$. In the pivot mechanism, the only participants who make or receive non-zero payments are ones who are pivotal.

Vickrey first introduced the pivot mechanism in a model where the decision x allocated a fixed quantity of a single divisible good. In the auction context, a bidder is not pivotal if he acquires a zero quantity. So the pivot mechanism in the Vickrey model is an auction in which losing bidders neither make nor receive payments.

2.2 Always Optimal and Weakly Dominant Strategies

In this section, we verify that the VCG rules do indeed ensure that it is always optimal for the participants to report truthfully, regardless of the reports made by others. We also demonstrate that reporting truthfully is often a *dominant strategy*, that is, it is the only strategy that is always optimal.

There are circumstances in which reporting truthfully, although always optimal for the VCG mechanism, is not a dominant strategy. For example, suppose that two parties are considering sharing the rental of a boat, which costs \$200. One party values the rental either at \$300 or at \$0, and his reported value is restricted to lie in the set {\$0, \$300}. The other party's value is some amount between \$0 and \$150, and his report is restricted to lie in the interval [\$0, \$150]. In this example, the pivot mechanism prescribes that the boat is rented if and only if the first party's value is \$300, and in that case the first party pays \$200. The second party always pays \$0, and his report does not affect the outcome. Consequently, any report by the second party is always optimal, and any report of \$200 or more by the first party is always optimal when his value is at least \$200.

The preceding example is constructed so participants can sometimes predict that certain reports will be irrelevant. In less contrived examples, one expects that truthful reporting will be a dominant strategy.

We formalize these claims using the following definitions. Truthful reporting is an *always optimal* strategy if condition (i) below holds, and it is a *dominant strategy*[1] if, in addition, condition (ii) holds:

(i) for all t^{-i}, $t^i \in \arg\max_{\bar{t}^i}\{v^i(\hat{x}(X, N, \bar{t}^i, t^{-i}), t^i) - \hat{p}^i(X, N, \bar{t}^i, t^{-i})\}$.

(ii) if $\bar{t}^i \neq t^i$, then for some t^{-i}, $\bar{t}^i \notin \arg\max_{\bar{t}^i}\{v^i(\hat{x}(X, N, \bar{t}^i, t^{-i}), t^i) - \hat{p}^i(X, N, \bar{t}^i, t^{-i})\}$.

To rule out contrived examples like the boat rental example, we will use the following condition:

All reports are potentially pivotal: For all $i \in N$ and $t^i, \bar{t}^i \in \Theta^i$, there exists $t^{-i} \in \Theta^{-i}$ such that $\sum_{j \in N} v^j(\hat{x}(X, N, \bar{t}^i, t^{-i}), t^j) < V(X, N, \vec{t})$.

This condition asserts that for any false report $\bar{t}^i$ by bidder i, there is some type profile t^{-i} of the other participants such that the false

[1] A strategy for a player in a normal form game is *dominant* if (1) it is a best reply to every opposing strategy profile and (2) there is no other strategy with the same property. The definition in the text specializes this definition to the direct revelation games we are studying.

report leads the mechanism to choose a suboptimal outcome. Whe this condition holds, no participant can be sure that a false report i harmless.

Theorem 2.1. In any VCG mechanism, truthful reporting is an *alway optimal* strategy. If all reports are potentially pivotal, then truthful re porting is a *dominant* strategy.

Proof. To show that truthful reporting is always optimal, fix the pro file $\vec{t}$ of actual types. When bidder i reports type $\bar{t}^i$, the decision cho sen is $\hat{x}(X, N, \bar{t}^i, t^{-i})$. So, given the formula for i's payment, his pay off is $\Pi^i(\bar{t}^i|\vec{t}) = v^i(\hat{x}(X, N, \bar{t}^i, t^{-i}), t^i) - \hat{p}^i(X, N, \bar{t}^i, t^{-i}) - h^i(t^{-i})$. Usin (2.4), the gain that i enjoys from the deviation is therefore

$$\begin{aligned}
&\Pi^i(\bar{t}^i|\vec{t}) - \Pi^i(t^i|\vec{t}) \\
&= [v^i(\hat{x}(X, N, \bar{t}^i, t^{-i}), t^i) - \hat{p}^i(X, N, \bar{t}^i, t^{-i}) - h^i(t^{-i})] \\
&\quad - [v^i(\hat{x}(X, N, \vec{t}), t^i) - \hat{p}^i(X, N, \vec{t}) - h^i(t^{-i})] \\
&= \sum_{j \in N} v^j(\hat{x}(X, N, \bar{t}^i, t^{-i}), t^j) - \sum_{j \in N} v^j(\hat{x}(X, N, \vec{t}), t^j) \\
&= \sum_{j \in N} v^j(\hat{x}(X, N, \bar{t}^i, t^{-i}), t^j) - V(X, N, \vec{t}) \leq 0.
\end{aligned}$$

This proves that truthful reporting is always optimal.

By the assumption that all reports are potentially pivotal, for all $\bar{t}^i \neq t^i$ there exists t^{-i} such that

$$\begin{aligned}
&\Pi^i(\bar{t}^i|\vec{t}) - \Pi^i(t^i|\vec{t}) \\
&= \sum_{j \in N} v^j(\hat{x}(X, N, \bar{t}^i, t^{-i}), t^j) - V(X, N, \vec{t}) < 0.
\end{aligned}$$

Hence, by definition, truthful reporting is a dominant strategy. ■

The formal proof implements the following simple intuitive argu ment. The VCG payments are defined so that i's report cannot affect the total payoff of the other participants. If i reports truthfully, the mech anism maximizes the total actual payoff. If i reports falsely in any way that changes the decision, then the change in total payoff must be neg ative and must be equal to the change in i's own payoff. So reporting truthfully is optimal. Moreover, if every false report is sometimes pivotal then it is sometimes suboptimal, so it is dominated by reporting truth fully.

The most widely known example of a pivot mechanism is the *second-price* auction. In the private-values auction model, a bidder's value for any decision depends only on the goods the bidder acquires, and not on the goods acquired by the other bidders: $v^i(x, t^i) = v^i(x^i, t^i)$, where $x^i = 1$ if the bidder acquires the good and $x^i = 0$ otherwise. The value of not acquiring the good is normalized to zero: $v^i(0, t^i) = 0$. Let us simply write v^i for $v^i(1, t^i)$.

Since losing bidders are not pivotal (because their presence does not affect the allocation), they pay zero in the pivot mechanism. According to (2.4), the price a winning bidder pays in this mechanism is equal to the difference between two numbers. The first number is the maximum total value to the other participants, including the seller, when i does not participate in the auction, which is $\max_{j\neq i} v^j$. The second number is the total value to the other bidders when i wins, which is zero. Thus, when bidder i wins, he pays $\max_{j\neq i} v^j$, which is equal to the second highest bid. For this reason, the pivot mechanism for the one good case is called the *second-price auction*.

Vickrey originally introduced the second-price auction as a model of ascending auctions, such as those now commonly used at internet auction sites. To develop the connection, we take special notice of the fact that auction sites like eBay and Amazon Auction encourage bidders to use a *proxy bidder* facility. The bidder tells the proxy a maximum price that it is willing to pay – its *maximum bid*. The proxy keeps this information secret and bids on the bidder's behalf in the ascending auction. Whenever it does not have the high bid, it raises the bid by one increment, provided that does not exceed the specified maximum bid. If every bidder were to use a proxy, then the result would be that the bidder who has specified the highest maximum price acquires the item and pays a price (approximately) equal to the second highest such price. If we replace the phrase "maximum price" with "bid price," then this is precisely the same rule that describes the outcome of a Vickrey auction for a single good. In the language of game theory, the English auction with proxy bidders and the second-price auction are *strategically equivalent*: there is a one-to-one mapping between the strategy sets such that corresponding strategy profiles lead to identical outcomes.[2]

[2] This theoretical account fairly describes Amazon Auction, but the rules are slightly different at eBay. eBay uses a fixed ending time after which no more bids are accepted. The ordering and timing of bid submissions can be relevant in an eBay auction. Indeed, "sniping" (waiting until the last few seconds to bid) is a common and viable strategy at eBay, but is almost totally absent at Amazon Auction, where an auction cannot end until there have been no new bids for ten minutes. See Ockenfels and Roth (2002).

We will henceforth use the term *Vickrey auction* to refer to the pivot mechanism in auction environments. By inspection of (2.4), the price paid by any participant $i \neq 0$ is equal to the loss imposed on other participants by adjusting the decision to account for i's values. This price is always non-negative. In contrast, prices paid in the more general VCG mechanism can be negative if h^i is sometimes negative. The possibility of negative payments to some participants raises a question about whether the sum of the payments to participants $i \neq 0$ is positive, negative, or zero.

2.3 Balancing the Budget

In public goods applications, the designer may want to ensure that the total payments to and from the participants *excluding the mechanism operator* add up to zero. This is called *balancing the budget*. If the mechanism designer is a public authority, this means that the authority runs neither a surplus nor a deficit on this project. In such cases, the mechanism designer typically has no independent value for the decision, so we formulate the model with $N = \{1, \ldots, n\}$, excluding the designer from the set of participants.

Definition. A direct mechanism (x, p) satisfies *budget balance* if for all finite Θ and all $\vec{t} \in \Theta$, the sum of all transfers is zero:

$$\sum_{i\in N} p^i(X, N, \vec{t}) = 0.$$

Summing the required payments reveals that the possibility of budget balance implies a restriction on the maximum value function, as follows:

$$0 = \sum_{i\in N} p^i(X, N, \vec{t}) = \sum_{i\in N} (p^i(X, N, \vec{t}) + h^i(t^{-i}))$$
$$= \sum_{i\in N} (V(X, N - i, t^{-i}) + h^i(t^{-i}))$$

$$- \sum_{i \in N} \sum_{j \in N-i} v^j(\hat{x}(X, N, \vec{t}\,), t^j)$$

$$= \sum_{i \in N} (V(X, N-i, t^{-i}) + h^i(t^{-i}))$$

$$- \sum_{i \in N} (V(X, N, \vec{t}\,) - v^i(\hat{x}(X, N, \vec{t}\,))$$

$$= (n-1)\left(\sum_{i \in N} f^i(t^{-i}) - V(X, N, \vec{t}\,) \right), \tag{2.5}$$

where

$$f^i(t^{-i}) = \frac{V(X, N-i, t^{-i}) + R^i(t^{-i})}{n-1}. \tag{2.6}$$

So a necessary condition for budget balance is that there exist functions f^i such that for all $\vec{t}$,

$$V(X, N, \vec{t}\,) = \sum_{i \in N} f^i(t^{-i}). \tag{2.7}$$

Holmstrom (1977) has observed that the same condition is actually necessary and sufficient for the existence of a budget-balancing VCG mechanism.

Theorem 2.2. There exists a VCG mechanism that satisfies budget balance if and only if there exist functions f^i such that (2.7) holds for all $\vec{t}$.

Proof. The necessity of (2.7) was established above. For sufficiency, given the functions f^i, take $h^i(t^{-i}) = (n-1) f^i(t^{-i}) - V(X, N-i, t^{-i})$ and observe that this implies (2.6) and hence (2.5). ■

To establish that the form (2.7) is restrictive, we use a simple two-player auction example with $N = \{1, 2\}$, a formulation that excludes the mechanism designer from the set of participants. There is a single good to be allocated, whose values to participants 1 and 2 are $v^1 \in \{1, 3\}$ and $v^2 \in \{2, 4\}$, respectively. There exists no way to represent $\max(v^1, v^2)$ as a sum $f^1(v^2) + f^2(v^1)$, so there can be no VCG mechanism in this setting that satisfies budget balance. To verify that directly, we tabulate the payments:

Participants' VCG Payments for the Four Value Profiles

	(1, 2)	(3, 4)	(1, 4)	(3, 2)
Participant 1	$h^1(2)$	$h^1(4)$	$h^1(4)$	$2 + h^1(2)$
Participant 2	$1 + h^2(1)$	$3 + h^2(3)$	$1 + h^2(1)$	$h^2(3)$
Total	$1 + h^1(2) + h^2(1)$	$3 + h^1(4) + h^2(3)$	$1 + h^1(4) + h^2(1)$	$2 + h^1(2) + h^2(3)$

Notice that, for any choice of h^1 and h^2, the sum of the total payments in the first two columns minus the corresponding sum in the last two columns is 1. Consequently, there is no choice of h^1 and h^2 such that all the column totals are zero: no balanced-budget VCG mechanism exists.

Theorem 2.2 still allows that there are some environments in which the VCG mechanism does always balance the budget. An important class of these are the ones in which the mechanism designer is treated as a participant who has just one possible type. In that case, the maximum value depends only on t^{-0} and so satisfies (2.7); indeed, $V(X, N, \vec{t}\,) \equiv f^0(t^{-0})$ for all $\vec{t}$. A VCG mechanism that works in this case specifies the pivot mechanism payments for all participants except participant 0 and balances the budget by having participant 0 receive the net proceeds of the mechanism. In situations where the mechanism designer is a regulator, a committee, or another entity with decision authority, the designer is frequently not allowed to receive or make payments from or to those over whom it has authority. Such restrictions might be imposed, for example, to prevent corruption in the system. In such cases, the budget-balance condition arises naturally and imposes restrictions on what can be implemented.

2.4 Uniqueness

Can another mechanism besides the VCG mechanism implement efficient decisions with dominant strategies? The answer depends on additional assumptions about the environment. For example, if there is a buyer whose value lies in the set {0, 10} and a seller whose cost of supplying a good is 5, then the following direct mechanism implements an efficient outcome in dominant strategies. In the mechanism, each party must report a value from its set of possible types. The seller has no choice but to report a cost of 5. If the buyer reports a value of 10, trade occurs at a price of 8; otherwise, there is no trade and no transfers occur. By inspection, it is a dominant strategy for both sides to report truthfully, and the outcome is always efficient. A VCG mechanism that makes no transfers when there is no trade is a pivot mechanism, and the pivot mechanism in this case sets a price of 5. It follows that the suggested mechanism is not a VCG mechanism.

The preceding example relied on the discrete nature of the type space. According to the next theorem, when the type space is smoothly connected, only the VCG mechanisms can implement efficient outcomes in dominant strategies.

Theorem 2.3. Suppose that for each i, $\Theta^i = [0, 1]$ (or simply that Θ^i is smoothly path connected[3]) and that for each decision outcome x, $v^i(x, t^i)$ is differentiable in its second argument. Then any efficient, incentive-compatible direct mechanism is a VCG mechanism. ■

The version of theorem 2.3 stated here was first proved by Holmstrom (1979), generalizing earlier work by Green and Laffont (1977), who had employed more restrictive assumptions about the type space. We postpone the proof to the next chapter, which contains several other closely related analyses.

2.5 Disadvantages of the Vickrey Auction

Despite its attractive features, the Vickrey auction has important disadvantages that make it unsuitable for most applications. In this section, we illustrate these disadvantages. We give a more detailed analysis of certain of the disadvantages in chapter 8, where the Vickrey design is pitted against certain leading alternatives.

The disadvantages of the Vickrey auction are divided into three kinds: practical disadvantages, monotonicity problems, and merger-investment disadvantages.

2.5.1 Practical Disadvantages

In this subsection, we discuss certain practical difficulties of implementing a Vickrey auction on account of factors that are omitted from the formal model.

[3] A set Θ is *smoothly path connected* if for every two points $\theta, \theta' \in \Theta$ there is a differentiable function $f : [0, 1] \to \Theta$ such that $f(0) = \theta$ and $f(1) = \theta'$.

One such problem is that a Vickrey auction can severely tax bidders' computational abilities. For example, consider a Vickrey auction to sell twenty spectrum licenses. In principle, each bidder must submit bids on every combination of licenses he might win, but there are more than one million such combinations. If the bidders must incur even a small cost to determine a value for each distinct combination of licenses, then the cost of running the Vickrey auction makes it impracticable. For some applications, this cost is not too onerous. For example, if the licenses are sufficiently similar, then a bidder might simply specify a value for each different number of licenses, or might adjust that for differences in the licenses. At least for the general case, allowing bids on all packages imposes costs that are too high for a realistic design.

A second practical problem is that real bidders often face serious budget limitations, which the Vickrey design does not take into account. In the presence of such constraints, a bidder in a Vickrey auction may have no always optimal strategy. For example, consider an auction with two identical goods and a bidder with values of 20 for one unit of a good and 40 for the package, but with a total budget of 10. This bidder has no always optimal strategy in the Vickrey auction. If there are credit restrictions or large penalties for default, then bids exceeding the bidder's budget can be ignored. If bidder 1's sole competitor bids 10 for one unit and 19 for two, his best reply is to bid 10 for one unit (and 10 for two units as well). However, if the competitor bids 9 for one unit only, then the best reply is to bid 0 for one unit and 10 for two units.

A third practical problem is that the Vickrey design may force the winning bidder to reveal too much information. A bidder might fear that its value information could be leaked, disadvantaging it in subsequent negotiations with the auctioneer or other buyers or suppliers (Rothkopf, Teisberg, and Kahn (1990)).

2.5.2 Monotonicity Problems

A different set of disadvantages of the Vickrey auction arises from the fact that payments are determined by a non-monotonic function of the bidders' values. We illustrate the problems that raises with a series of examples, borrowed from Ausubel and Milgrom (2002). A formal analysis that identifies the set of auction environments in which these disadvantages are relevant is presented in chapter 8, as part of a comparison of the advantages of several multi-object auction designs.

Here, we provide a series of examples illustrating the monotonicity problems that the Vickrey auction can suffer. In the Vickrey auction, (1) *adding* bidders can *reduce* equilibrium revenues, (2) revenues can be zero even when competition is ample, (3) even losing bidders can have profitable joint deviations in which they *increase* their bids in concert to win items while creating *lower* prices for themselves, and (4) bidders can profitably use shill bidders, intentionally increasing competition in order to generate lower prices.

Consider a Vickrey auction of two identical spectrum licenses. Bidders 1 and 2 are new entrants, which each need two licenses to establish a business of economic scale. Bidder 1 is willing to pay up to $1 billion for the pair of licenses, and bidder 2 is willing to pay up to $900 million. If these are the only bidders in the auction, then the auction is effectively a second-price auction for the pair of licenses. Bidder 1 will acquire the two licenses for a price of $900 million.

Now, suppose instead that there are two additional bidders. Bidders 3 and 4 are both incumbent wireless operators. Each seeks just a single additional license to expand the capacity of its network. Suppose each incumbent is willing to pay up to $1 billion for a single license. If the Vickrey auction is used and all bidders play their dominant strategies, then the two incumbents will acquire the licenses. Because the licenses are given to those who value them the most, this outcome is efficient and results in a total value of $2 billion.

One might expect that increasing the number of bidders and their maximum total value for the pair of licenses would increase the seller's revenue, but that is not the case: the total price paid by the winning bidders is *zero*. To see why, let us compute the price paid by bidder 3. According to (2.4), this price is the opportunity cost to the other bidders of the license that bidder 3 wins. More specifically, it is the maximum value of the two licenses to the other three bidders, which is $1 billion, minus the maximum value of a single license to those bidders, which is also $1 billion. The difference of zero is bidder 3's price and bidder 4's price is determined in the same way.

Notice that the declining revenue problem vanishes if the first two bidders regard the licenses as substitutes. For example, suppose that instead of bidding only $1 billion for two licenses, bidder 1 is also willing to pay $500 million for one license, and similarly bidder 2 is willing to pay $450 million for one license. Then bidders 3 and 4 must each pay $500 million for a license, and the seller's revenue climbs from $900 million to $1 billion.

The next two variations exploit the feature of the Vickrey auction that, when goods are not substitutes, prices may decrease as the bids increase or the set of bidders expands.

First, we modify the preceding example. As before, bidders 1 and 2 each want only a pair of licenses and are willing to pay $1 billion or $900 million for the pair, respectively. In the modified example, however, each of the incumbents, bidders 3 and 4, has a value of $400 million for a single license. If the incumbents play their dominant strategies, they win no licenses and earn payoffs of zero. If, however, they act in concert, both raising their bids to $1 billion for a single license, then the prices are determined just as above, and the situation is the one we have already examined: bidders 3 and 4 win the two licenses for a total price of zero. Thus, the Vickrey auction provides opportunities and incentives for collusion among the low-value, losing bidders.

Next, we consider another variation. In this one, there are only three bidders, with the first two described just as above. In this variation, the third bidder is also a new entrant and also has value only for the pair of licenses, but its value is lower than that of the first two bidders. It is willing to pay just $800 million for the pair of licenses, compared to $900 million and $1 billion for the other two bidders. Still, the third bidder can win the licenses profitably by entering the auction with two identities, as bidders 3 and 4, and having 3 and 4 each bid $1 billion for a single license. The result, just as before, is that bidders 3 and 4 win, each acquiring a single license for a price of zero. Thus, by combining the tactics of shill bidding and loser collusion, a bidder in the Vickrey auction whose values are too low to be assigned any licenses at the efficient allocation can profitably win both licenses and force the seller to accept a zero price.

Standard auctions do not suffer the monotonicity problems plaguing the Vickrey auction. For example, if the seller simply takes sealed bids and awards licenses to the highest bidders at prices equal to the winning bids, then none of the monotonicity problems occur: Adding bids and bidders cannot reduce prices; introducing shill bids cannot reduce anyone's price, and losers cannot become winners except by paying higher prices.

These monotonicity problems are significant practical defects. In section 2.5.3 below, we reexamine these examples to see whether they are in some sense exceptional, that is, whether they are unlikely to arise in practice. We find that, to the contrary, monotonicity problems can only be ruled out in cases where goods are likely to be substitutes, which is a small subset of the possible cases.[4]

2.5.3 The Merger–Investment Disadvantage

The Vickrey auction also suffers another important disadvantage, distinct from those described above. This one arises even when the auctioneer's objective is efficiency rather than revenue, and when shill bidding and collusion are impossible. The problem is that the Vickrey design can distort the bidders' investment and merger incentives *ex ante* (before the auction),[5] leading to inefficiency.[6]

To illustrate, we return to the first example of the previous section, in which bidders 1 and 2 value only the pair of licenses and have values of $1 billion and $900 million, respectively. Suppose that, before the auction, bidders 3 and 4 could merge and, by coordinating, increase the total value of the licenses by 25% from $2 billion to $2.5 billion. Even though such a merger would increase the maximum total value, the parties would not profit by merging. Recall that the unmerged firms paid a total of zero and enjoyed net profits of $2 billion. The merged firm, however, would pay $1 billion in a Vickrey auction, leaving it a net profit of just $1.5 billion.

In this example, the Vickrey auction discourages a merger by reducing the joint profits of the merging parties. Thus, even by the standard of efficiency, the Vickrey mechanism can have significant disadvantages.

[4] In an unpublished result, Daniel Lehmann has shown that with more than two items, the restriction that items must be substitutes fails generically. That is, treating the valuation functions as a vector, for any valuation v where goods are substitutes, almost every valuation in any neighborhood of v fails to satisfy the substitutes condition.

[5] Several authors have developed analyses based on the observation that there are no such distortions for single item auctions. With the set of bidders fixed, because any bidder's profit is equal to his contribution to social surplus, the bidder has correct incentives for any investments that affect only his own values. The same applies to bidders' decisions about how much information to acquire about their own values (Bergemann and Valimaki (2002)).

[6] Economists typically emphasize market power issues when analyzing mergers, and those issues are excluded entirely from this analysis. As discussed earlier, the term "efficiency" as used in mechanism design theory is narrower than the economic idea of Pareto optimality, because here it takes into account only the interests of the mechanism participants.

In analyzing merger incentives, as in studying collusion and shill bidding, whether the assets being auctioned are substitutes proves important. In the Vickrey auction, if the bidders regard the goods as substitutes, then winners generally can reduce their prices by merging. Thus, Vickrey auctions tend to favor mergers when goods are substitutes. For example, suppose that there are four bidders for three items. Each of the first three bidders has a value of 2 for a single item and the fourth bidder has a value of 1. The Vickrey outcome is that the three high-value bidders acquire single items for a price of 1. If the first two bidders merge, the allocation of goods is the same: the merged bidder gets two units and bidder 3 gets one unit. Bidder 3's price is unchanged – it pays a price of 1 for its unit – but the merged bidder pays a total of 1 for its *two* units, so its average price is $\frac{1}{2}$ per unit. This price reduction is typical for the case when goods are substitutes.

If the government is to auction assets to an industry in which it wishes to promote competition or encourage entry, e.g. electrical power generation, it may properly view with suspicion rules that promote mergers and favor larger bidders.

As our examples have shown, however, Vickrey auctions do not always promote mergers. In our telecommunications auction example, we found that merged firms may pay relatively high prices and may even find it profitable to use shills to divide demand between two smaller bidders. If shills are impossible, then the Vickrey auction may discourage profitable and welfare-enhancing mergers. Taken together, the various examples establish that Vickrey auctions can be too favorable to mergers or too discouraging.

2.6 Conclusion

The Vickrey–Clarke–Groves theory provides important insights into what mechanism design can achieve. In the class of environments with quasi-linear preferences, the VCG mechanisms provide every participant with a dominant strategy, which is to reveal his type truthfully. When bidders do report honestly, the decision selected is the total-value-maximizing one. Moreover, the VCG mechanisms are the *only* mechanisms that exhibit these two properties without restrictions on the possible set of values.

Offsetting these advantages of the VCG mechanisms are certain problems. Using the VCG mechanism to decide how much of a public good to produce may prevent balancing the budget. Budget balance presents no obstacle to using the VCG mechanism to conduct an auction, however, for the auctioneer is quite happy to pocket any surplus that the mechanism generates.

Besides the budget balance problem, the Vickrey auction suffers a variety of other drawbacks. Some of these are practical, associated with the complexity of the auction, its inability to accommodate budget constraints, and the information it demands from the bidders. Another set of drawbacks are the *monotonicity problems*, which include the possibility that increased competition can lead to reduced seller revenues, that revenues can be very low or zero even when competition is substantial, that losing bidders may have profitable ways to collude, and that a single bidder can sometimes benefit by pretending to be several independent bidders. The third set of drawbacks concern distortions in merger and related investment decisions.

We return to the monotonicity problems in chapter 8, where we will find that they are potentially present in a wide range of environments. They are reliably absent only if all bidders regard all the goods being sold as substitutes. In chapter 8, we will identify an alternative mechanism that matches the advantages of the Vickrey design when goods are substitutes but avoids some of the disadvantages.

In the chapters between, the VCG mechanism plays a very different role – as a benchmark for assessing the performance of alternative mechanisms.

REFERENCES

Ausubel, Lawrence and Paul Milgrom (2002). "Ascending Auctions with Package Bidding." *Frontiers of Theoretical Economics* **1**(1): Article 1.

Bergemann, Dirk and Juuso Valimaki (2002). "Information Acquisition and Efficient Mechanism Design." *Econometrica* **70**(3): 1007–1033.

Clarke, E.H. (1971). "Multipart Pricing of Public Goods." *Public Choice* **XI**: 17–33.

Green, Jerry and Jean-Jacques Laffont (1977). "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods." *Econometrica* **45**: 427–438.

Groves, Theodore (1973). "Incentives in Teams." *Econometrica* **61**: 617–631.

Holmstrom, Bengt (1977). *On Incentives and Control in Organizations*: Doctoral thesis, Stanford University.

Holmstrom, Bengt (1979). "Groves Schemes on Restricted Domains." *Econometrica* **47**: 1137–1144.

Ockenfels, Axel and Alvin E. Roth (2002). "Last Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet." *American Economic Review*: **92**(4): 1093–1103.

Rothkopf, Michael, Thomas Teisberg, and Edward Kahn (1990). "Why Are Vickrey Auctions Rare?" *Journal of Political Economy* **98**: 94–109.

Vickrey, William (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* **XVI**: 8–37.