CHAPTER 1

Introduction

1.1 Decentralized allocation mechanisms

Suppose each of you in the class has an *endowment* of some apples and some pears. Some of you like apples more than pears; some like pears more than apples. Each of you can rank all baskets of apples and pears by a *preference relation* such that for all baskets you can tell whether you prefer a particular basket to another or you are *indifferent* between them. For example, you may prefer to have 7 apples and 6 pears to having 6 apples and 7 pears, so that $\{7a, 6p\} \succ \{6a, 7p\}$, where $\succ$ stands for "prefers." In turn, you may be indifferent between having a basket of 7 apples and 6 pears and a basket of 4 apples and 11 pears, so that $\{7a, 6p\} \backsim \{4a, 11p\}$. Your preferences can be also represented by a *utility function* $U_i(a, p)$, increasing in both arguments. If $\{7a, 6p\} \succ \{6a, 7p\}$, then $U(7a, 6p) > U(6a, 7p)$. If $\{7a, 6p\} \backsim \{4a, 11p\}$, then $U(7a, 6p) = U(4a, 11p)$. Utility is just some number **you**, i, attach to the value of a particular basket.

Suppose now that each of you wants to *maximize* the utility you derive from consuming apples and pears by exchanging them with others. Once the exchange is completed, but not before, you eat the apples and the pears you have.

How can this exchange process be organized? Suppose we do it like this:

(1) Each of you puts your apples and your pears on the table. When this is done, a bell rings and each of you takes all the apples and pears you want and leaves the room with them.

What will happen? Obviously, a fight. What will be the final *allocation*? Bullies will take more and some black-and-blue marks will be inflicted.

Is this *allocation mechanism* a "market"? It does not quite look like one, since people grab, using force, rather than exchange voluntarily. There is nothing to safeguard your endowments and nothing to protect your allocation: only your big muscles and swift feet. But notice that this is a *decentralized mechanism*: Everyone acts independently. Moreover, the allocation process is completely free: You do what you want.

The outcome is bad. Some people will get injured. If it is possible to generate the same allocation of apples and pears without fighting, some people would be better off (without bruises) without anyone being worse off. Hence, if people care about bruises, this outcome is *inefficient*: Someone can be better off without anyone being worse off. Moreover, strong people may get a lot; weak people nothing, which somehow does not seem right. Even our most vague moral intuitions tell us that the allocation should be related to something like *effort* or *need*, not just the good *luck* of being physically strong. It should be *equitable*. Hence, this is not a good way to allocate.

(2) Suppose then we do it in a different way. All the apple owners can exchange their apples for pieces of paper on which pear owners write the number of pears they will give in return for each apple tomorrow. Hence, i gives j an apple and j gives i a piece of paper that says "p": the pear *price* j is willing to pay tomorrow.

Let us ask again what will happen. What will be the final allocation? The obvious answer is that nothing will be exchanged. The final allocation will be the same as the original endowments. The reason is that if the exchange were to occur, it would not be in the interest of j tomorrow to fulfill her promise. Hence, the promise is not *credible*: j may well write some number "p" on a piece of paper, but she will not deliver on her promise. In turn, if i knows that j will not fulfill her promise, i will not alienate his apples.

Note that this mechanism looks more like a market. An exchange could now occur. But something is missing, namely, *enforcement institutions*. The fact that the promises concern tomorrow does not matter. Suppose that they were for one hour from now, one minute from now, or even one second from now. In the end, the exchange of apples for pears may be almost simultaneous: The question still remains who will lift his hand off the apples and pears first.

Is the final allocation a good one? Let us stipulate that i has $\{3a, 2p\}$ but would like more $\{1a, 3p\}$ while j has $\{1a, 2p\}$ but prefers $\{3a, 1p\}$. Hence, if i could give j $2a$ in exchange for $1p$, they would both be happier. Because the initial endowments can be improved on and because no exchange occurs under this allocation mechanism, the final allocation is inefficient. Hence, the mechanism "exchange apples for promises of pears" is free and decentralized but it is bad in the sense of leading to an inefficient allocation.

(3) Yet another way. Suppose that there is someone who can and will, can **and** will, force the owners of pears to deliver them to the sellers of apples in quantities written on the pieces of paper, so that the pieces of paper are as good as real pears. Moreover, to make exchanges easier, one can also write pieces of papers promising apples in exchange for pears.

What will the outcome be? First, we need to answer the following question: Will the exchange process ever stop? There are many students in this class, each with some initial endowment. Now i sells $2a$ to j for $1p$. But k comes along and offers to buy each of i's pears for a price of $3a$. Suppose that, given the number of apples i has, he is willing to give up a pear for 3 apples. Then i will sell her pears to k. But, and so on. As you see, the stopping question is not a trivial one. But you know from your microeconomics course – and we will cover this material remedially in the next chapter – that this exchange will come to a stop. It will reach some allocation such that no one will want to trade anymore, an *equilibrium*.

What will be true of this equilibrium? First, all the potential gains from trade will be exhausted: No one will be able to benefit from continuing the exchange. Second, no one can be better off without someone else being worse off: The only way i could increase her utility is by someone else losing his. Finally, if we were to take a vote whether to alter the equilibrium allocation, then at least one person would vote against changing it: This allocation would not be defeated under unanimity rule. *These three conditions are equivalent, and they fully describe efficiency in the sense of Pareto or Pareto optimality.*

Note that we have smuggled in a number of assumptions. One is that no one eats during the exchange process. If you were to consume before the equilibrium allocation is reached, *out of equilibrium*, then some of you would be consuming apples or pears for which you would have paid too little or too much. You may have consumed an apple for which you paid $2p$ while there may have been someone who would have sold you an apple for $1p$, so that you could have consumed more apples and pears. We assumed that trading is costless, that is, that it has no *transaction costs*.

Second, we assumed that everyone knows everything and everyone knows the same. Suppose that the pears are not all the same – some are rotten – and that only the sellers know their quality. Then you may suspect that if someone is willing to sell you a pear for a low apple price, this may be because she knows that the pear is of poor quality. You may be unwilling to buy pears at all, so that there may be no market for pears (Ackerlof 1970).

Suppose for the moment that these assumptions hold. We know that the equilibrium allocation will be Pareto efficient. Does it mean that this mechanism that generated this allocation is a good one? In the example we discussed thus far the initial endowments and the equilibrium allocation were

$$i : 3a, 2p \Rightarrow 1a, 3p$$
$$j : 1a, 2p \Rightarrow 3a, 1p$$

But suppose that the initial endowments were different and the exchange led to a different allocation:

$$i : 2a, 2p \Rightarrow 0a, 3p$$
$$j : 2a, 2p \Rightarrow 4a, 1p$$

Can these equilibrium allocations be compared by the Pareto criterion? Because the utility functions increase in both arguments, i is better off in the original than in the modified world: She has the same number of pears but one apple more. But j is better off in the new world: He has the same number of pears as in the original one but one more apple. Hence, j is better off but i is worse off. The Pareto criterion cannot be used to compare these two allocations: We cannot judge which is more efficient by this criterion. Each equilibrium allocation is *Pareto superior* to the corresponding endowment allocation: Otherwise there would have been no exchange. But each initial endowment leads to a different equilibrium allocation, and neither of them is Pareto superior to the other.

Hence, if we want to evaluate these states of the world, we need some other criteria. We could, for example, ask which of these worlds is better in terms of total societal consumption or total utility. In the first case, we would use as the criterion a function $W(C_i, C_j)$, where C stands for consumption, such that

$$W(C_i, C_j) = C_i + C_j.$$

According to this criterion, an allocation that makes this sum larger is better than one that makes it smaller. Because we are not producing anything and because nothing is lost in exchange, total consumption is the same in the two worlds, so this criterion does not work. But even if total consumption associated with different allocations was different, this criterion would be vulnerable to the criticism that the value of consumption to different individuals may be different.

So we are left with total utility, the classical utilitarian way to think of *social welfare*. This criterion is a function $W(U_i, U_j)$ such that

$$W(U_i, U_j) = U_i(C_i) + U_j(C_j).$$

In terms of our example, the original world is better than the modified one if

$$U_i(1a, 3p) + U_j(3a, 1p) > U_i(0a, 3p) + U_j(4a, 1p),$$

or

$$U_i(1a, 3p) - U_i(0a, 3p) > U_j(4a, 1p) - U_j(3a, 1p).$$

As you see, to make this evaluation, we have to compare the utilities of i and j, which is hard, if not impossible, to do. The Us are numbers individuals attach and i may attach 3 to his difference but may as well attach 102, so how are we to tell if i's difference is greater than j's? Moreover, because i prefers $\{1a, 3p\}$ to $\{0a, 3p\}$, if we asked i what her value of the difference is, she would say "1000": She would reveal her preferences strategically. Hence, this is not a promising route either.

How about comparing these allocations by the criterion of *equality*? Is this a good criterion? Note that in our example the world in which the equilibrium allocation would be perfectly equal would not be a good world: We already know that people would want to trade away from equal endowments. *Outcome-egalitarianism* does have its supporters. But even if we are egalitarians, we must ask: Equality of what? (Sen 1992).

"Utility" is an appealing answer: People should be able to be equally happy with their consumption baskets. If they do not envy other people's allocations, that is, if given their utility functions both i and j prefer their own allocation to the allocation of the other person, then these allocations are at least *fair*. But we have seen that as an operational criterion utility is hard to implement.

Perhaps what we should equalize is *opportunity*. Suppose now that what you consume is produced. Specifically, the output (you can think of it as total value of apples and pears measured in equilibrium prices) is produced by two factors of production: luck (L) and effort (E). We can write the production function generally as

$$Y = F(L, E)$$

increasing in both arguments. Suppose that the specific production function is simply

$$Y = L + E$$

Now, say that i has 8 lucks and exerts 2 efforts, while j exerts the same amount of effort but has no luck. Their output will be:

$$i : 8 + 2 = 10$$
$$j : 0 + 2 = 2$$

Note that their opportunities are very different: i can get 8 units of consumption doing nothing, while j must rely completely on her effort. We could equalize the opportunity for consumption by taking 4 lucks away from i and giving them to j. They will both have an opportunity of 4. Obviously, they may exert different amounts of effort and end up with different consumption. But their rewards will be, in Dworkin's (1981a and b) language, "ambition-" rather than "endowment-" sensitive.

But what is "luck" and what is "effort"? If I watch TV the whole day rather than work, is it because I am lazy or because I have a deficiency of laboramine in my brain? And how do luck and effort combine in production? Suppose that you observe Y but not L and E separately. How to equalize opportunity now? (The problem is not hopeless, see Roemer 1996.)

Is opportunity a good *equalisandum*? Suppose we equalized opportunity, so that everyone has 4 lucks and i turns out to be a bum, exerting no effort, and starves having only 4 to consume. Should we accept this allocation or should we ensure that everyone has some basic consumption basket? You may say with Giddens (1998), "no rights without responsibilities": i had the same chance to produce the basic basket as everyone else and if she decided to squander her opportunities it is just tough. But suppose that the reason she did not exert any effort was that she fell ill. Again, you may say that she should have insured herself against this possibility; to anticipate what follows, you may even say that she should have been *forced* to insure herself. But suppose she was not forced and did not buy insurance: A hiker climbs a mountain, breaks a leg, and is lying on a ledge, dying from the cold. A helicopter televises his agony. He could have bought rescue insurance but he did not. Should he be allowed to just die or should he be saved?

Finally, you may want to evaluate allocations by a *maximin* criterion (Rawls 1971). Suppose that we rank all the baskets of some basic goods from the largest to the smallest. Then an allocation is better if it makes larger the smallest basket.

1.2 Centralized allocation mechanisms

There are obviously other ways we can allocate and other ways we can evaluate the allocations and the mechanisms that generate them. Let us consider a few more.

(4) Here is the most obvious: i is a *dictator* and she decides who gets what for everyone. What individuals get does not depend on their actions, just on the decision of the dictator, who can enforce it using physical coercion if need be. Will the resulting allocation be a good one? Is this a good allocation mechanism? Do not jump to conclusions. Whatever the dictator decides, the allocation will be Pareto efficient: Any other allocation would make someone worse off, namely, the dictator. But this may just show that the Pareto criterion is too weak. More interestingly, the dictator may be benevolent: She may just want to implement one of the criteria we discussed above.

One question arises immediately: Will a decentralized mechanism generate the same allocation as a benevolent and omniscient dictator? This question is often posed by economists as a **counterfactual** method to evaluate allocations: It makes sense to think that an allocation chosen by a benevolent, omniscient dictator is the best possible, so that any mechanism that deviates from the *command optimum* must be in some way deficient. This is, however, an excessively demanding standard: How would the dictator know everything? We can weaken it by assuming that the dictator knows no more and not less than individuals and ask about the *constrained command optimum.* For example, we can ask whether the allocation that results from the last decentralized mechanism we considered is *constrained Pareto efficient*, that is, whether this allocation could be Pareto improved on by a benevolent dictator who knows only what the trading individuals know.

Yet even this weaker criterion is still counterfactual. Individuals know things the dictator does not know: Most obviously what makes them happy but perhaps also how much they need to survive or how much luck they enjoy and how much effort they exert. If a benevolent dictator is to act in their interest, he must somehow elicit this information. But individuals will *reveal* their private information only if they have *incentives* to do so. This may mean that even constrained Pareto efficiency may be unattainable to the dictator.

Note that what is at stake in these distinctions is what we can realistically expect as the best possible. Many debates about the virtues and vices of different allocation mechanisms hinge on this question. We will return to such issues several times.

(5) Let us consider another centralized mechanism. Say the Central Authority confiscates the initial endowments of apples and pears, puts them in baskets as before, and runs a lottery to determine who gets which basket. Once the results of the lottery are known, the Central Authority distributes to individuals the baskets they won.

Before we consider how to evaluate this mechanism, let us pause on the notion of a *centralized mechanism*. In what sense is this mechanism centralized? Is it because we have a Central Authority? But in the example of decentralized exchange of commodities for commodities through pieces of paper, that we also had some kind of an authority that used or threatened coercion if people did not deliver on their pieces of paper. We have already seen that there can be no exchange, at least no generalized anonymous exchange, without some kind of an enforcement mechanism.

It is useful to ask first why we did not pose this question with regard to the dictator. The answer is that in the case of the dictator, it is obvious that one decision generates the entire allocation to all the individuals. The dictator decides how much to give to each. But so does the lottery. Lots are thrown into an urn, a handsome television actor, smiling broadly, reaches into the urn to allocate consumption baskets to everyone. Contrast this with exchange, where i decides whether to sell to j but their exchange does not affect the allocation to k. Exchange is a decentralized mechanism because the final allocation results from independent decisions of each agent; dictatorship and lottery are centralized mechanisms because one decision allocates to everyone.

Is lottery a good mechanism? It does have its virtues. One would want to say intuitively that it is fair, but we have seen that economists reserved this term for an allocation that does not produce envy, and lottery certainly does. It equalizes chances, but they are not quite the opportunities we discussed above, because all you can do with your basket is to consume it. It is obviously inefficient and outcome-inegalitarian. We tend to employ lotteries, I suspect, when we think that exchange is not an ethically defensible mechanism and have no obvious criteria by which to allocate (on these issues, see Elster 1992). For example, some countries use lotteries to allocate scarce medical resources, including body parts. We do not think they should be exchangeable for money and we cannot see other clear criteria. Hence, we leave things to luck.

(6) One more, I promise the last, allocation mechanism. Suppose we use majority vote. Say there are three individuals so that we can easily use the

Table 1.1

	Apples	**Pears**
i	1	3
j	3	1
k	2	2
Total	6	6

majority criterion. The mechanism works as follows. Any individual can propose any allocation $\{\mathbf{a}, \mathbf{p}\}$ where, note, the members of this allocation are column vectors of size 3, because there are three individuals. The only condition we will impose is that of a *balanced budget*, so that the sums allocated must equal the total endowments available. For example, say that individual i proposes the following allocation to all three.

The proposal of i is then paired against the status quo (say the initial endowments of each individual) and people cast their votes either for the status quo or for i's proposal. Whichever of the two alternatives wins is the new status quo and everyone, i included, can propose a different allocation, which will be paired against the status quo, and so on. Note that this mechanism is a centralized one: Everyone together decides how to allocate to each.

The first thing we need to do is to return to the stopping problem. Will this process ever end? Or will it continue ad infinitum? Because this is an issue shrouded in confusion, let us be careful about the question we are asking. We are not asking how much time it will take to reach the final decision about allocation. This was not the question about exchange either: We asked in fact whether there exists some allocation such that, **whenever** it is reached, exchange will stop. So let us ask the same question now: Is there some allocation such that, whenever it is reached, voting will stop? Suppose that there is some allocation that beats every other possible allocation by majority vote. Say the proposal of i, which we will call X, does it. If j proposes some other allocation, say Y, $X \succ_M Y$ (read $\succ_M$ as "beats by" or "is preferred to" under majority rule). If k proposes Z, $X \succ_M Z$, and the same is true for every other proposal. Then X is the *majority (or Condorcet) winner*. And if X is the majority winner, then the voting process will "stop" at X: Nothing will defeat it. The electorate will have chosen to allocate consumption according to X, given in the Table 1.1.

Before we raise a problem, let us consider the qualities of this allocation. Is it efficient? This is the same as to ask whether individuals would want to trade their allocations after the voting process had stopped. You will see that i's proposal gives i and j the same they were being allocated in our example of free exchange. Hence, if this were the allocation, trading would stop (assuming that k does not want to trade). Moreover, someone would vote against any other allocation. So it is efficient.

The problem is that voting would not stop at X. Even worse, there is no proposal that beats every other alternative by a simple majority rule. Just think: Say that in response to X, j proposes an allocation Y that gives less to i and more to j and k. Then the latter two will vote for Y over X, so that $Y \succ_M X$. But then k will propose an allocation Z that gives more to k and i than Y, so that these two will not vote for it and $Z \succ_M Y$. But i can offer her initial proposal again. This proposal will give i and j more than Z, $X \succ_M Z$. Hence, we have $X \succ_M Z \succ_M Y \succ_M X$. There is no proposal that beats every other proposal under pairwise majority rule. Majority winner does not exist, which means that this mechanism is not *decisive*: It fails to pick one from among all alternatives.[1]

As you see, this is a different kind of a deficiency than those we discussed above. The problem here is not that the equilibrium allocation is somehow undesirable but that there is no allocation that constitutes an equilibrium. We just do not know what to expect in this situation. Perhaps this only means that we did not describe the voting mechanism adequately, but this is for later.

1.3 Political-economic equilibria

Let us first summarize what we have done thus far. We discussed three decentralized mechanisms of allocation: "grab all you can," "exchange commodities for pieces of paper," and "exchange commodities for commodities via pieces of paper." Then we discussed three centralized mechanisms: dictatorship, lottery, and majority rule. With regard to each of them, we asked first which allocation they will generate. We discovered that in some cases this question has a relatively easy answer but in one case the answer may be

[1] The word "cycling" is banned from this book. Nothing "cycles" here. This paragraph does not describe how proposals will be made but only a property of the function that transforms individual preferences into a collective one. All that we have learned is that this function fails to pick a unique allocation.