Nova Scotia (3%)

Saskatchewan (3%)

Manitoba (4%)

Alberta (11%)

British Columbia (13%)

Quebec (23%)

Ontario (39%)

In a yes–no voting system, any collection of voters is called a *coalition*. A coalition is said to be *winning* if passage is guaranteed by yes votes from exactly the voters in that coalition. Coalitions that are not winning are called *losing*. Thus, every coalition is either winning or losing. In Example 1, the coalition made up of France, Germany, and Italy is a winning coalition, as is the coalition made up of France, Germany, Italy, and Belgium. Note that when one asserts that a collection of voters is a winning coalition, nothing is being said about *how* these players actually voted on a particular issue. One is simply saying that *if* these people voted for passage of some bill and the other players voted against passage of that bill, the bill would, in fact, pass.

For most yes–no voting systems, adding extra voters to a winning coalition again yields a winning coalition. Systems with this property are said to be *monotone*. For monotone systems, one can concentrate on the so-called *minimal winning coalitions*: those winning coalitions with the property that the deletion of one or more voters from the coalition yields a losing coalition. In our example above, France, Germany, and Italy make up a minimal winning coalition, while France, Germany, Italy, and Belgium do not. (In the European Economic Community, the minimal winning coalitions are precisely the ones with exactly twelve votes—see Exercise 2.)

2.3 WEIGHTED VOTING AND THE U.N. SECURITY COUNCIL

The four examples of yes–no voting systems in the last section suggest there are at least three distinct ways in which a yes–no voting system can be described:

1. One can specify the number of votes each player has and how many votes are needed for passage. This is what was done for the European Economic Community. More generally, if we start with a set of voters, then we can construct a yes–no voting system by assigning real number *weights* to the voters (allowing for what we might think of as either a fractional number of votes for some voter, or even a negative number of votes) and then set any real number q as the "quota." A coalition is then declared to be winning precisely when the sum of the weights of the voters who vote "yea" meets or exceeds the quota. (Even in the real world, quotas can be less than half the sum of the weights—see Exercise 1.)

2. One can explicitly list the winning coalitions, or, if the system is monotone, just the minimal winning coalitions. This is essentially what is done in the description of the U.S. federal system above, since the three clauses given there describe the three kinds of winning coalitions in the U.S. federal system. In fact, if one deletes the parenthetical clauses and the phrase "or more" from those descriptions, the result is a description of the three kinds of minimal winning coalitions in the U.S. federal system.

3. One can use some combination of the above two, with provisos that often involve veto power. Both the U.N. Security Council and the procedure to amend the Canadian Constitution are described in this way. Moreover, the description of the U.S. federal system in terms of the tie-breaking vote of the vice president, the presidential veto, and the Congressional override of this veto is another example of a description mixing weights with provisos and vetoes. (We say "weights" in this context since,

for example, one can describe majority support in the House by giving each member of the House weight one and setting the quota at 218, and two-thirds support from the House by setting the quota at 290.)

The following observation is extremely important for everything we shall do in this chapter: A given yes–no voting system typically can be described in more than one way.

For example, instead of using weights, we could have described the European Economic Community by listing the fourteen winning coalitions. A similar comment applies to the U.N. Security Council. In fact, *every* yes–no voting system can be described by simply listing the winning coalitions. Conversely, any collection of subsets of voters gives us a yes–no voting system, although most of the systems arrived at in this way would be of little interest.

These observations lead to the central definition of this chapter.

DEFINITION. A yes–no voting system is said to be a *weighted system* if it can be described by specifying real number weights for the voters and a real number quota—with no provisos or mention of veto power—such that a coalition is winning precisely when the sum of the weights of the voters in the coalition meets or exceeds the quota.

Example 1 (not surprising): The European Economic Community is a weighted voting system.

There is no doubt about this, since we described the system by explicitly producing the weights and the quota. Consider, however:

Example 2 (surprising): The U.N. Security Council is also a weighted voting system.

This is not obvious. Although our description of this voting system did involve weights (weight one for each of the fifteen members) and a quota (nine), it also involved the statement that each of the five permanent members has veto power.

So how does one show that a given yes–no voting system is, in fact, weighted? The answer is that one must find (that is, *produce*) weights for each of the voters and produce a quota q so that the winning coalitions are precisely the ones with weight at least q. The difficulty with doing this for the U.N. Security Council is that we must somehow assign extra weight (but how much?) to the permanent members in such a way that this weight advantage alone builds in the veto effect.

Here is the intuition behind our method of finding a set of weights and a quota q that will prove the U.N. Security Council is a weighted voting system. Since all of the nonpermanent members clearly have the same influence, we will begin by assigning them all the same weight, and we will take this weight to be 1. The five permanent members also have the same influence and so we will (temporarily) assign them all the same unknown weight x. Now, exactly what must x and q satisfy for this to work?

Consider a coalition made up of all ten nonpermanent members together with any four permanent members. This has weight $4x + 10$, and must be losing since the one permanent member not in the coalition has veto power. Thus $4x + 10 < q$. On the other hand, the five permanent members together with any four nonpermanent members is a winning coalition, and so $q \leq 5x+4$. Putting these two inequalities together yields $4x + 10 < 5x + 4$, and so $6 < x$.

The previous two paragraphs suggest that in our pursuit of weights and quota, we try weight 1 for the nonpermanent members and weight 7 for the permanent members. Our inequalities

$$4x + 10 < q \quad \text{and} \quad q \leq 5x + 4$$

now imply that $38 < q \leq 39$. Thus, we certainly want to try $q = 39$.

It will now be relatively easy to demonstrate that these weights and quota do, indeed, show that the U.N. Security Council is a weighted voting system. (We should also note that to prove a system is weighted, it suffices to produce the weights and quota and show they "work." It is not necessary to explain how you found them.) The argument we seek runs as follows.

PROPOSITION. *The U.N. Security Council is a weighted system.*

PROOF. Assign weight 7 to each permanent member and weight 1 to each nonpermanent member. Let the quota be 39. We must now show that each winning coalition in the U.N. Security Council has weight at least 39, and that each losing coalition has weight at most 38.

A winning coalition in the U.N. Security Council must contain all five permanent members (a total weight of 35) and at least four nonpermanent members (an additional weight of 4). Hence, any winning coalition meets or exceeds the quota of 39. A losing coalition, on the other hand, either omits a permanent member, and thus has weight at most

$$(7 \times 4) + (1 \times 10) = 28 + 10 = 38,$$

or contains at most three nonpermanent members, and thus has weight at most

$$(7 \times 5) + (1 \times 3) = 35 + 3 = 38.$$

Hence, any losing coalition falls short of the quota of 39. This completes the proof.

It turns out that if one alters any weighted voting system by giving one or more of the voters veto power, the resulting yes–no voting system is again a weighted voting system. See Exercises 5 and 6 at the end of the chapter.

A natural response to what we have done so far is to conjecture that *every* yes–no voting system is weighted. Perhaps, for example, even for the U.S. federal system, we can do something as clever as what we just did for the U.N. Security Council to find weights and a quota that work. Alas, this turns out not to be the case, as we will see in the next section.

2.4 SWAP ROBUSTNESS AND THE NONWEIGHTEDNESS OF THE FEDERAL SYSTEM

The U.S. federal system, it turns out, is not a weighted voting system. But how does one prove that a system is *not* weighted? Surely we cannot simply say that we tried our very hardest to find weights and a quota and nothing that we tried appeared to work. Asserting that a system is not weighted is saying no one will *ever* find weights and a quota that describe the system. Moreover, we cannot check all possible choices of weights and quota since there are infinitely many such choices.

Here is one answer. To prove that the U.S. federal system is not weighted, it suffices to find a property that we can prove

1. holds for every weighted voting system, and
2. does not hold for the U.S. federal system.

One such property is given in the following definition:

DEFINITION. A yes–no voting system is said to be *swap robust* if a one-for-one exchange of players (a "swap") between two winning coalitions X and Y leaves at least one of the two coalitions winning. One of the players in the swap must belong to X but not Y, and the other must belong to Y but not X.

Thus, to prove that a system is swap robust, we must start with two *arbitrary* winning coalitions X and Y, and an *arbitrary* player x that is in X but not in Y, and an *arbitrary* player y that is in Y but not in X. We then let X' and Y' be the result of exchanging x and y. (Thus, x is now in Y' but not X', while y is in X' but not Y'.) We must then show that either X' or Y' is winning. To illustrate this, consider the following.

PROPOSITION. *Every weighted voting system is swap robust.*

PROOF. Assume we have a weighted voting system and two arbitrary winning coalitions X and Y with X containing at least one voter x not in Y and Y containing at least one voter y not in X. Suppose now that voter x from the winning coalition X is exchanged for voter y from the winning coalition Y to yield X' and Y' as above. If x and y have the same weight then both X' and Y' are winning, since X' weighs the same as X and Y' weighs the same as Y. If, on the other hand, x is heavier than y, it then follows that Y' weighs strictly more than Y, since Y' was obtained by deleting y and adding the heavier x. Thus the weight of Y' certainly exceeds the quota, and thus Y' is winning as desired. (In this latter case, X' may or may not be winning.) If y is heavier than x, then the argument is analogous to what we just gave. This completes the proof.

Our goal now is to show that a particular yes–no voting system—the U.S. federal system—is *not* swap robust. To do this we must *produce*